



US009455975B2

(12) **United States Patent**
Brooker et al.

(10) **Patent No.:** **US 9,455,975 B2**
(45) **Date of Patent:** **Sep. 27, 2016**

(54) **TECHNIQUES FOR MANAGING CREDENTIALS IN A DISTRIBUTED COMPUTING ENVIRONMENT**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Amazon Technologies, Inc.**, Reno, NV (US)

7,058,807 B2 * 6/2006 Grawrock G06F 21/57
713/172
8,752,047 B2 * 6/2014 Banga G06F 9/5077
718/1

(72) Inventors: **Marc J. Brooker**, Cape Town (ZA);
Mark Joseph Cavage, Seattle, WA (US);
David Brown, Cape Town (ZA);
Kevin Ross O'Neill, Seattle, WA (US);
Eric Jason Brandwine, Haymarket, VA (US);
Christopher Richard Jacques de Kadt, Cape Town (ZA)

2005/0251808 A1 11/2005 Gbadegesin et al.
2007/0107048 A1 * 5/2007 Halls G06F 21/33
726/4
2007/0167151 A1 * 7/2007 Zinn H04W 4/12
455/411
2010/0266128 A1 * 10/2010 Asokan H04L 63/0853
380/278
2011/0002341 A1 * 1/2011 Damola et al. 370/401
2011/0197065 A1 * 8/2011 Stauth G06F 21/45
713/170
2012/0072972 A1 * 3/2012 Christiansen H04L 63/0815
726/5

(73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

2012/0110576 A1 5/2012 Lin et al.
2012/0185913 A1 * 7/2012 Martinez et al. 726/1
2013/0013738 A1 * 1/2013 Astete G06F 9/45533
709/217

* cited by examiner

(21) Appl. No.: **14/204,124**

(22) Filed: **Mar. 11, 2014**

Primary Examiner — Saleh Najjar

Assistant Examiner — Louis Teng

(65) **Prior Publication Data**

US 2014/0196130 A1 Jul. 10, 2014

(74) *Attorney, Agent, or Firm* — Davis Wright Tremaine LLP

Related U.S. Application Data

(63) Continuation of application No. 12/981,231, filed on Dec. 29, 2010, now Pat. No. 8,683,560.

(51) **Int. Cl.**

H04L 29/06 (2006.01)

G06F 21/44 (2013.01)

H04L 9/32 (2006.01)

(52) **U.S. Cl.**

CPC **H04L 63/08** (2013.01); **G06F 21/44** (2013.01); **H04L 9/3247** (2013.01)

(58) **Field of Classification Search**

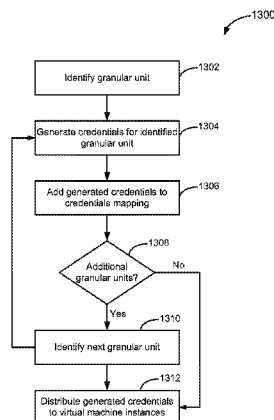
CPC H04L 67/10; H04L 63/10; H04L 63/20; G06F 21/53; G06F 2221/2149

See application file for complete search history.

(57) **ABSTRACT**

Systems and methods for managing credentials distribute the credentials to subsets of a set of collectively managed computing resources. The collectively managed computing resources may include one or more virtual machine instances. The credentials distributed to the computing resources may be used by the computing resources to perform one or more actions. Actions may include performing one or more functions in connection with configuration, management, and/or operation of the one or more resources, and/or access of other computing resources. The ability to use credentials may be changed based at least in part on the occurrence of one or more events.

20 Claims, 16 Drawing Sheets



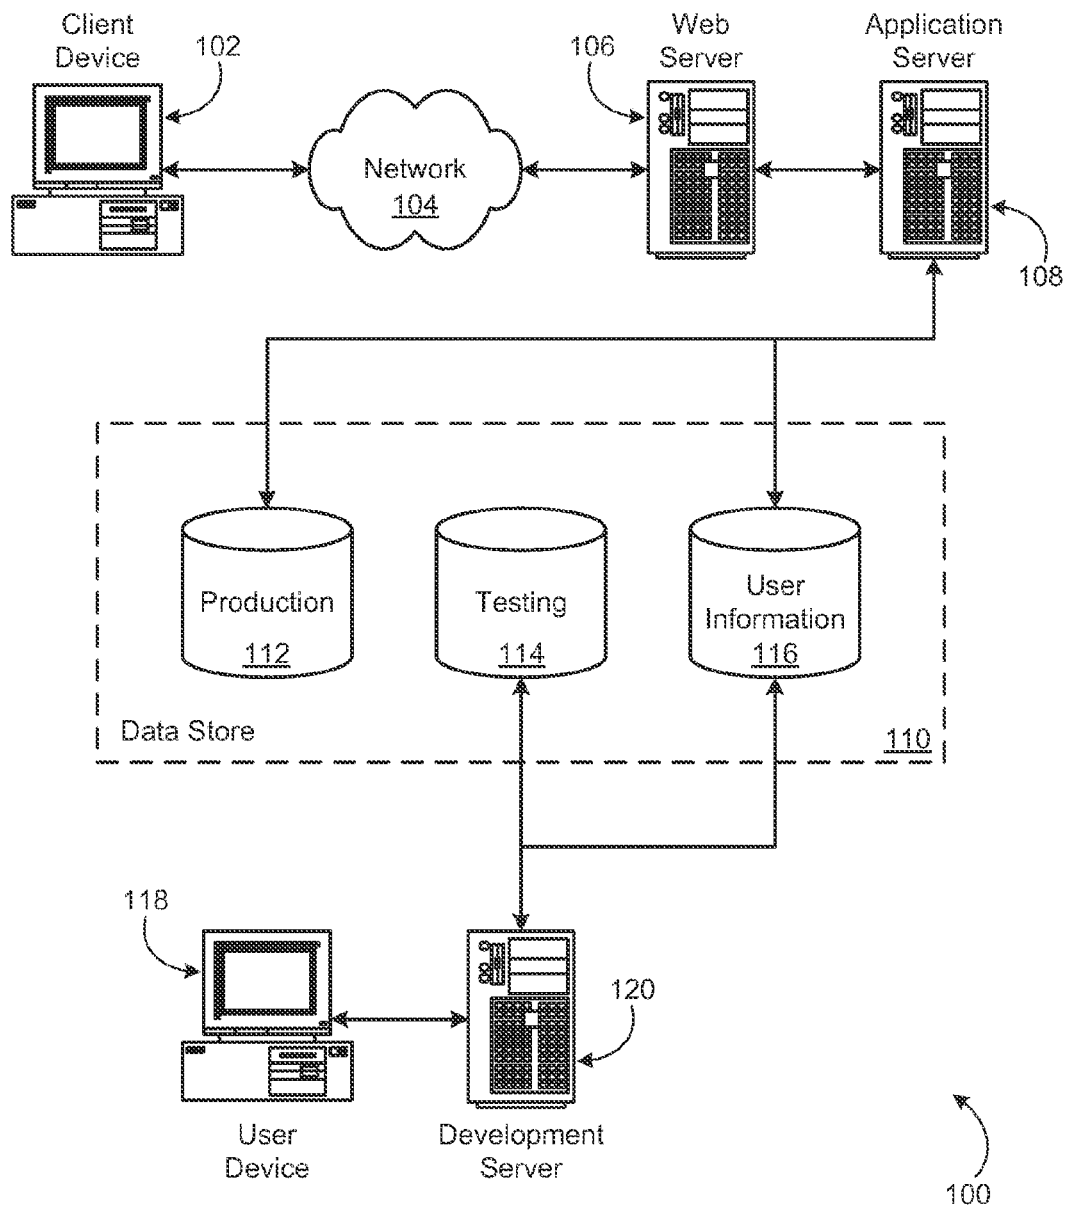


Figure 1




Figure 2

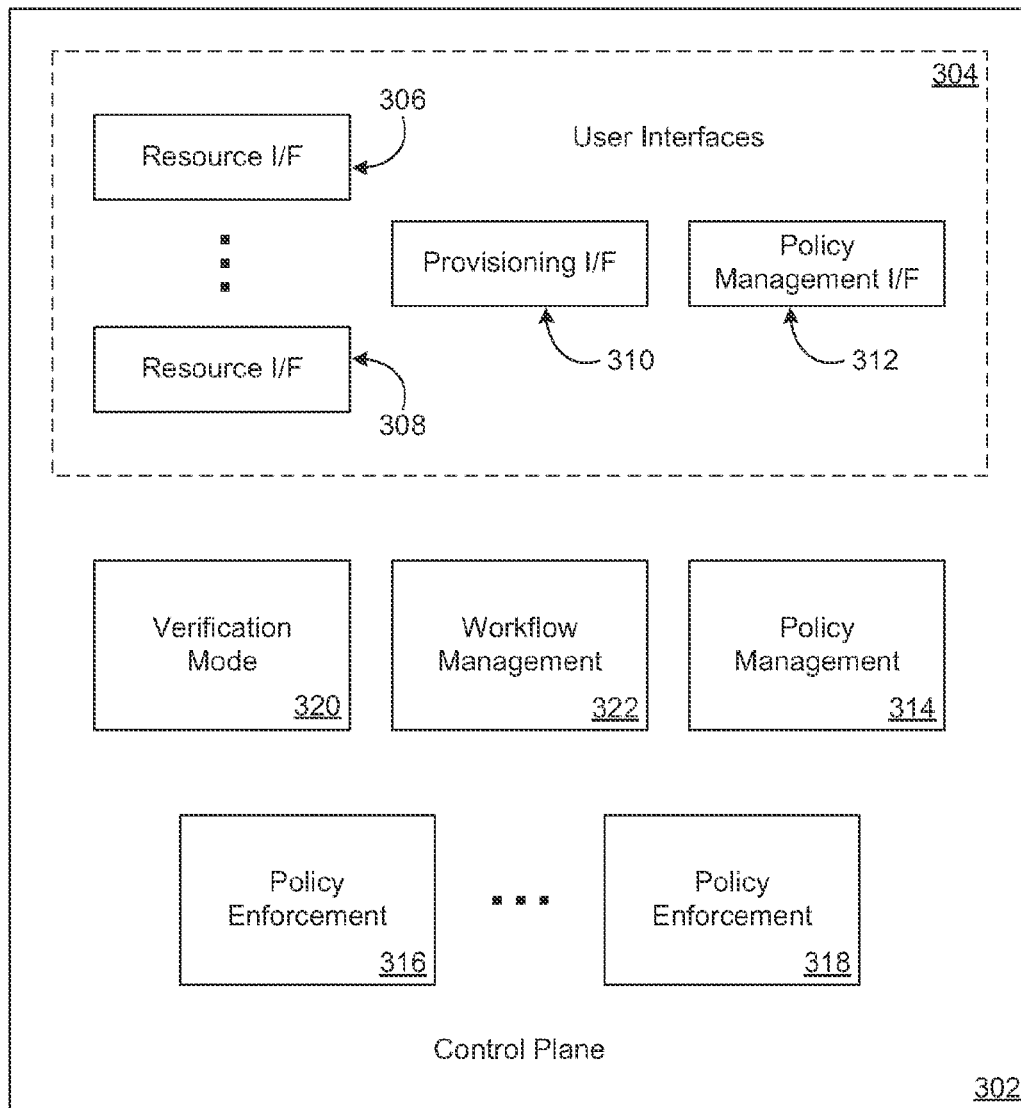


Figure 3

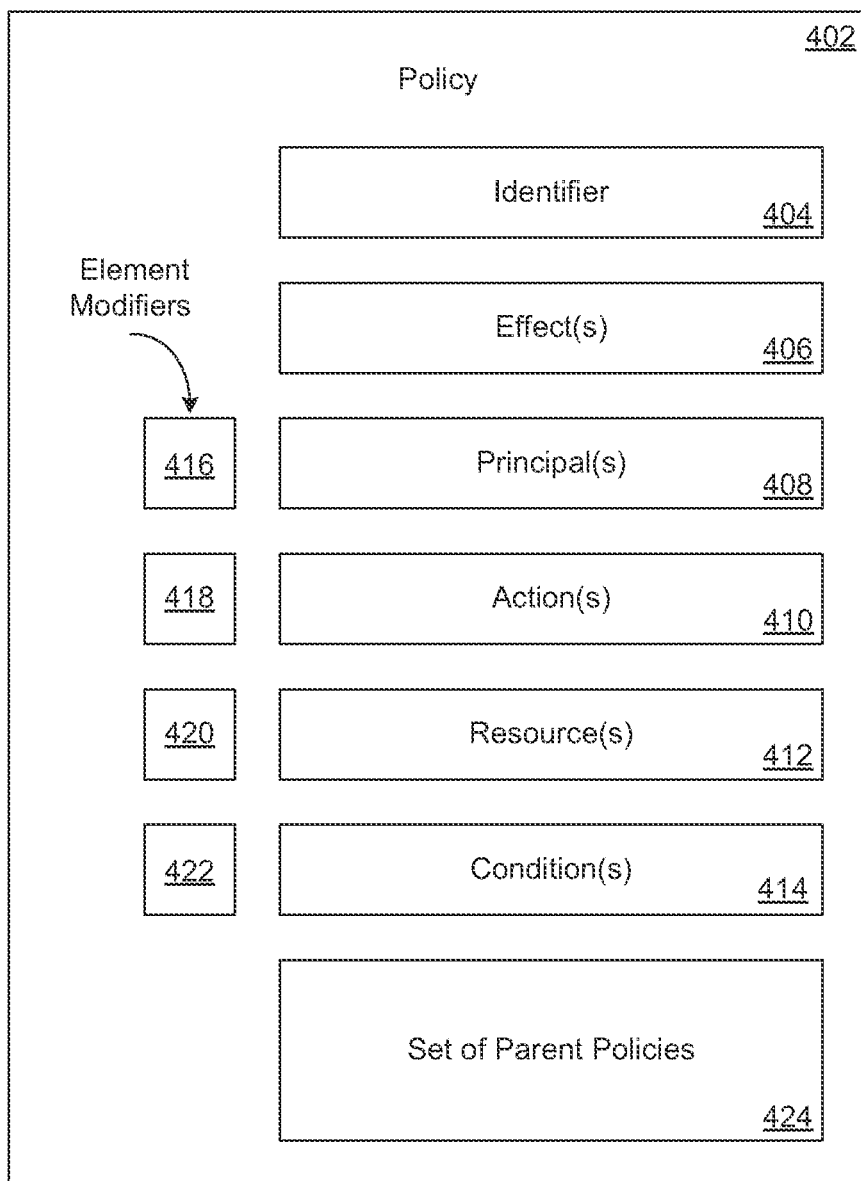


Figure 4

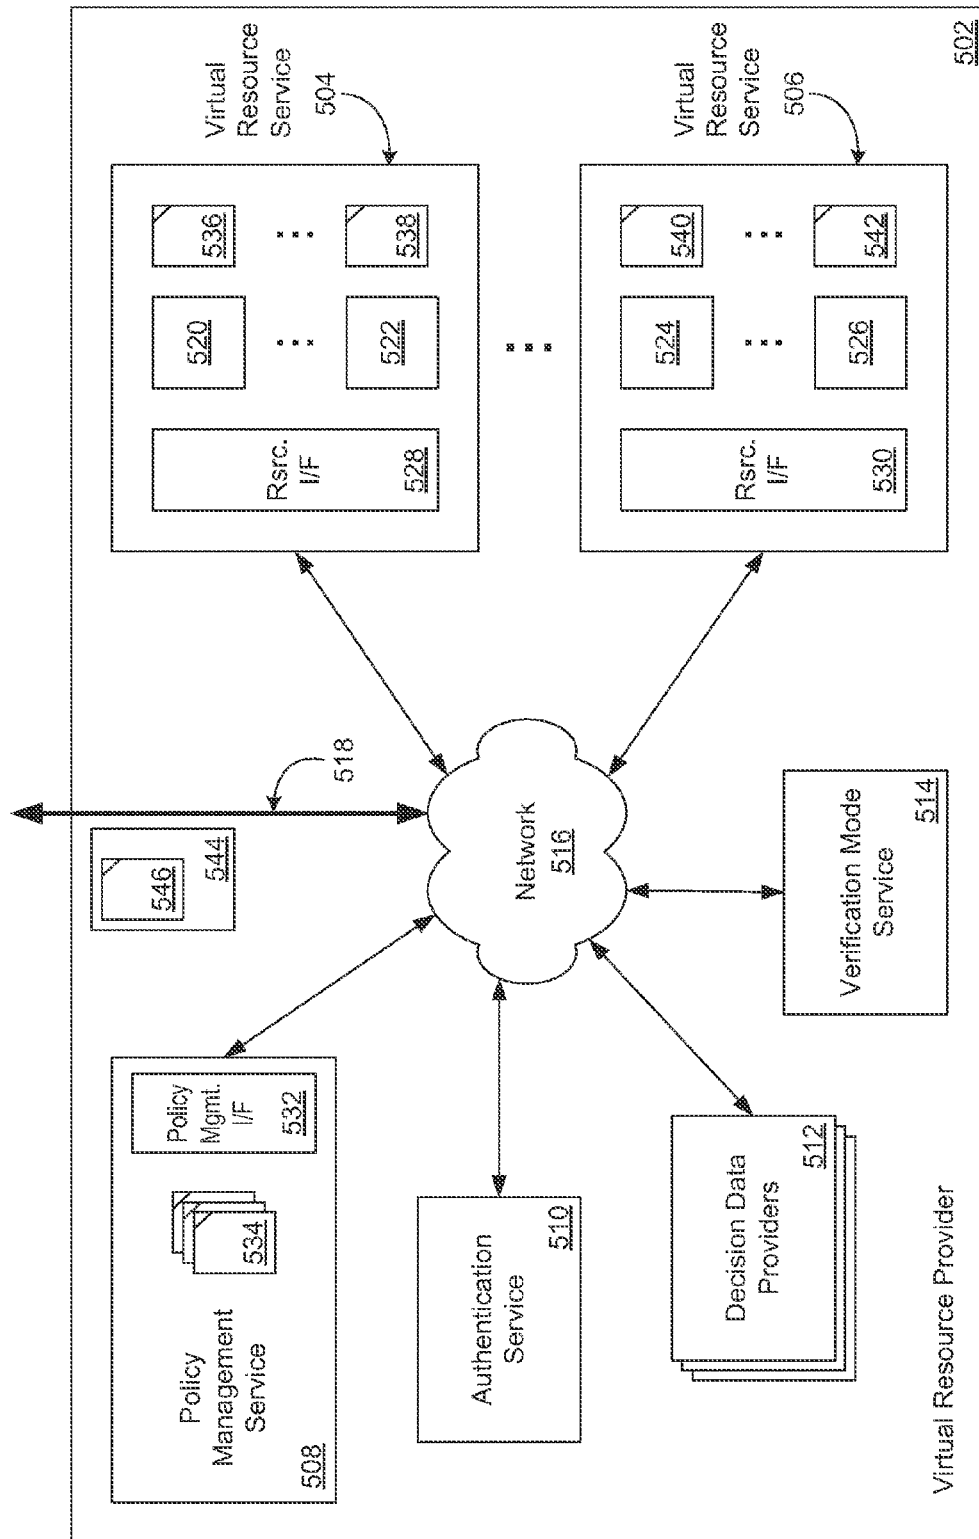


Figure 5

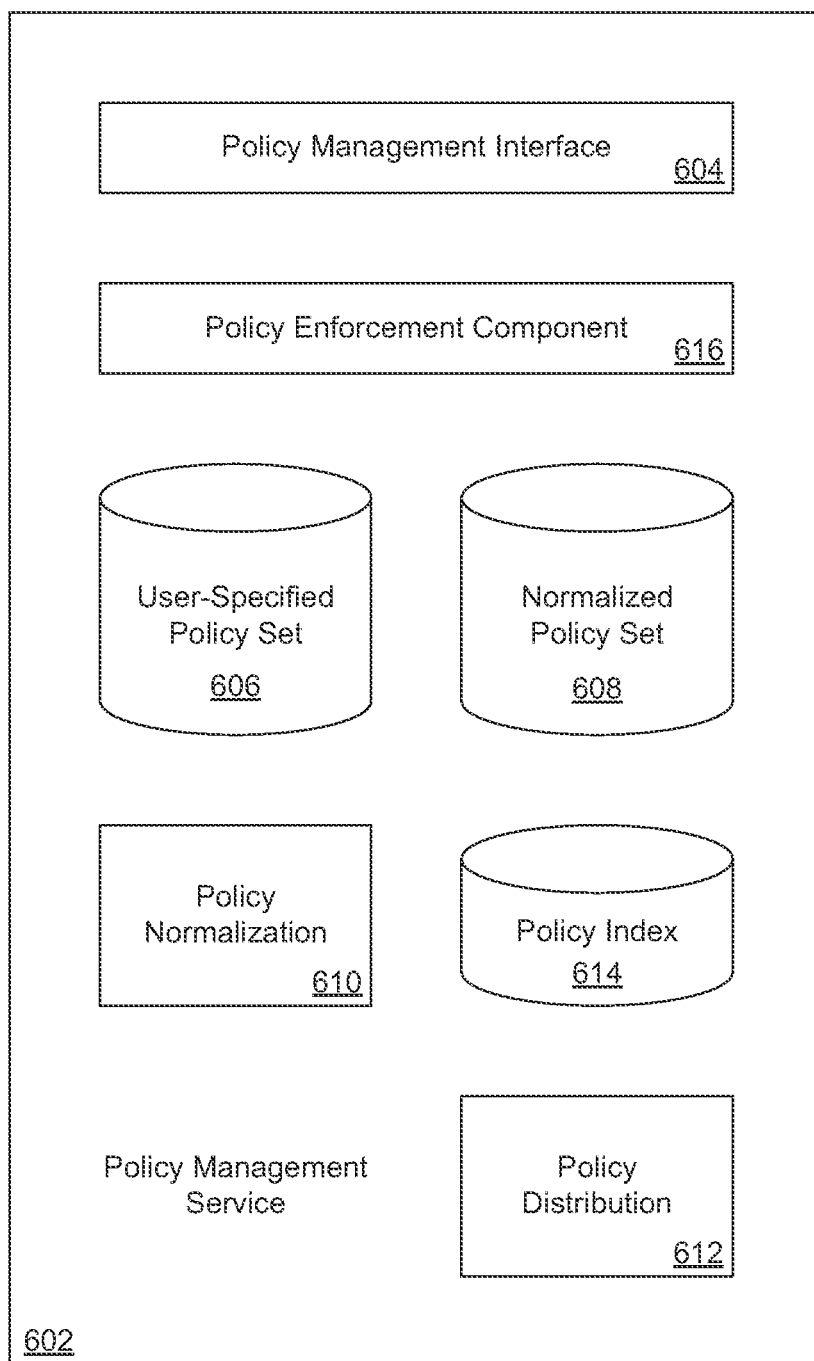


Figure 6

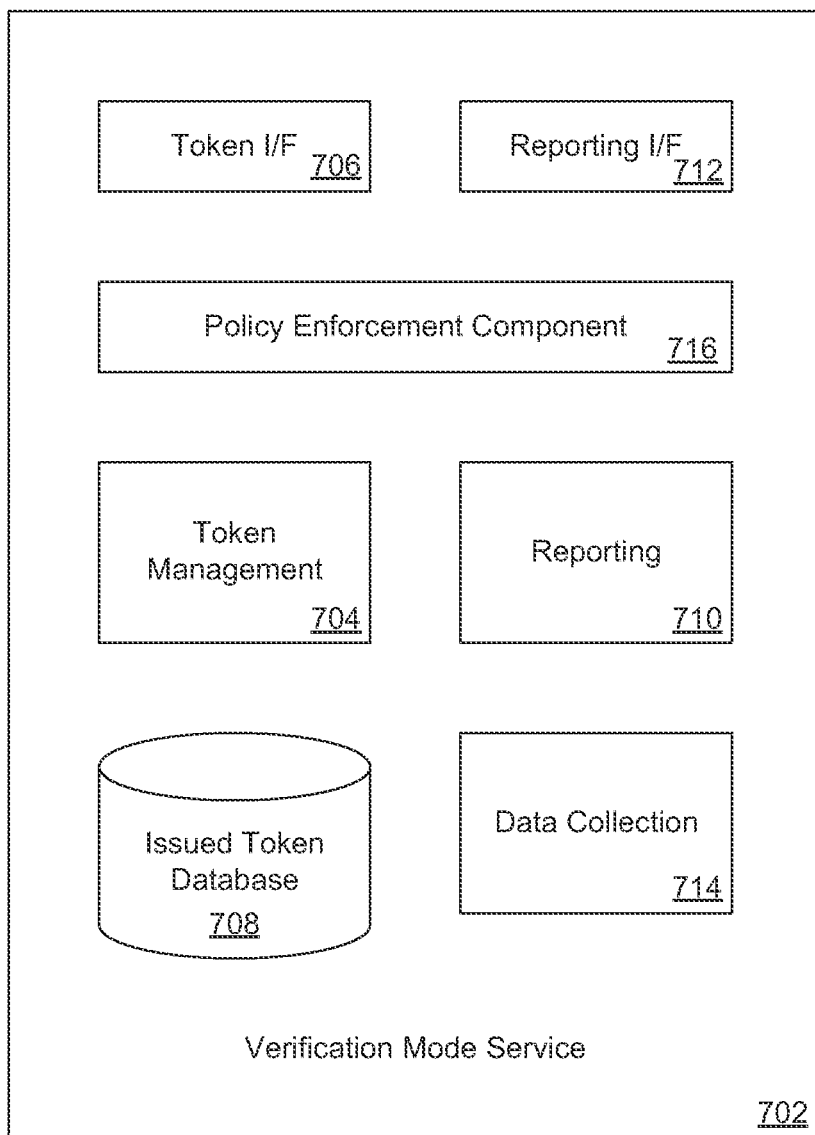


Figure 7

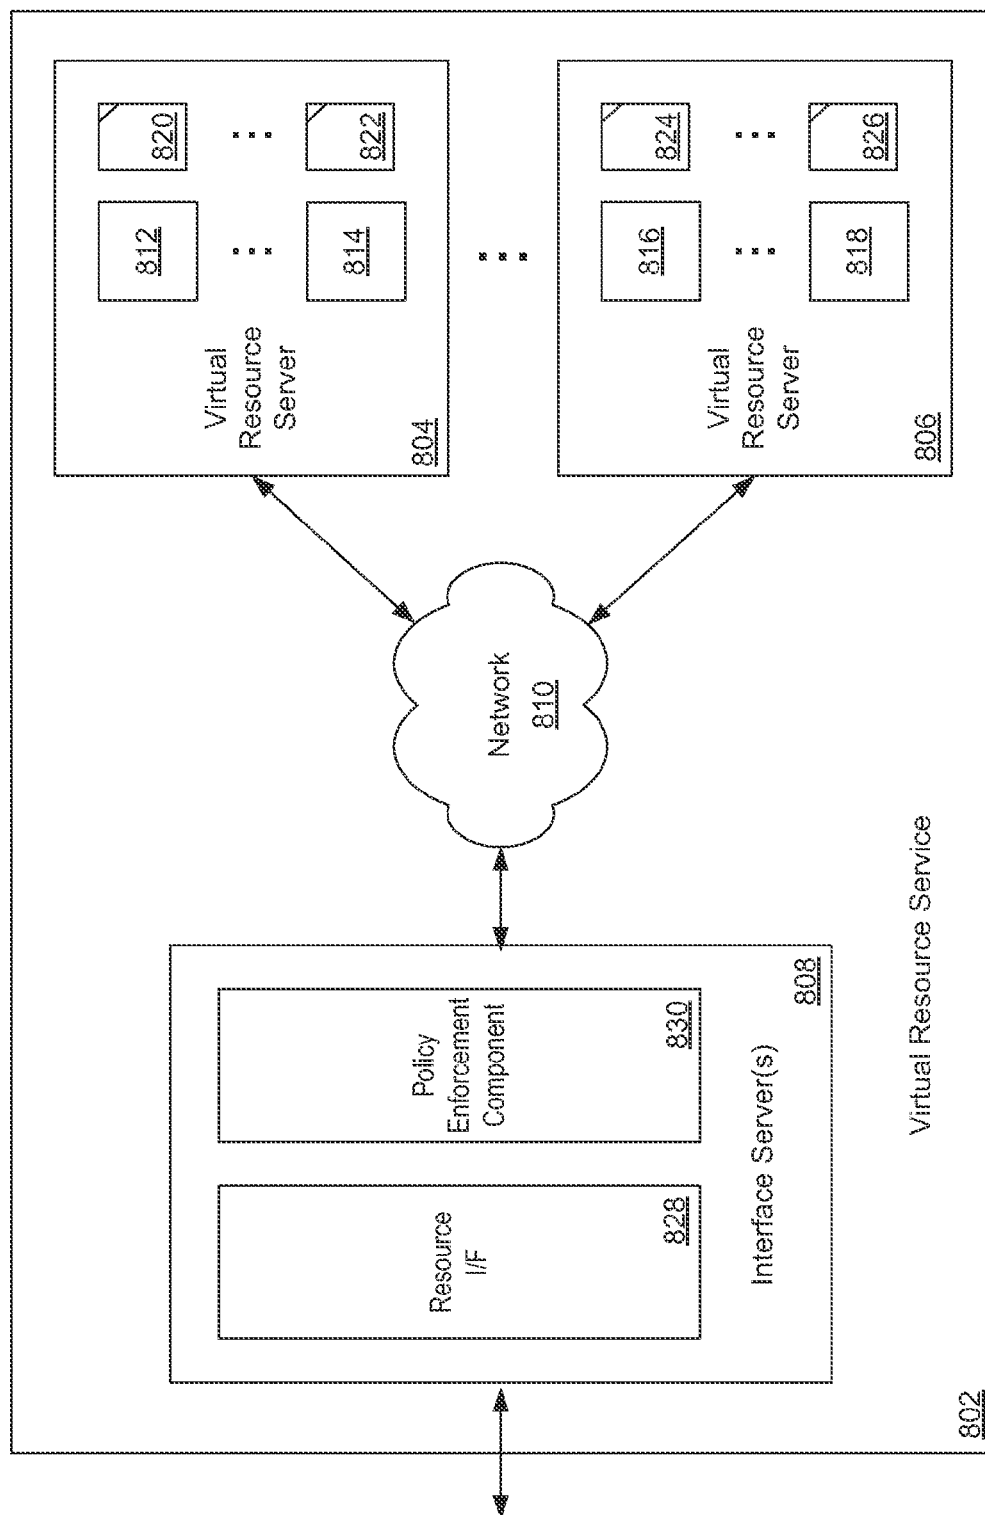


Figure 8

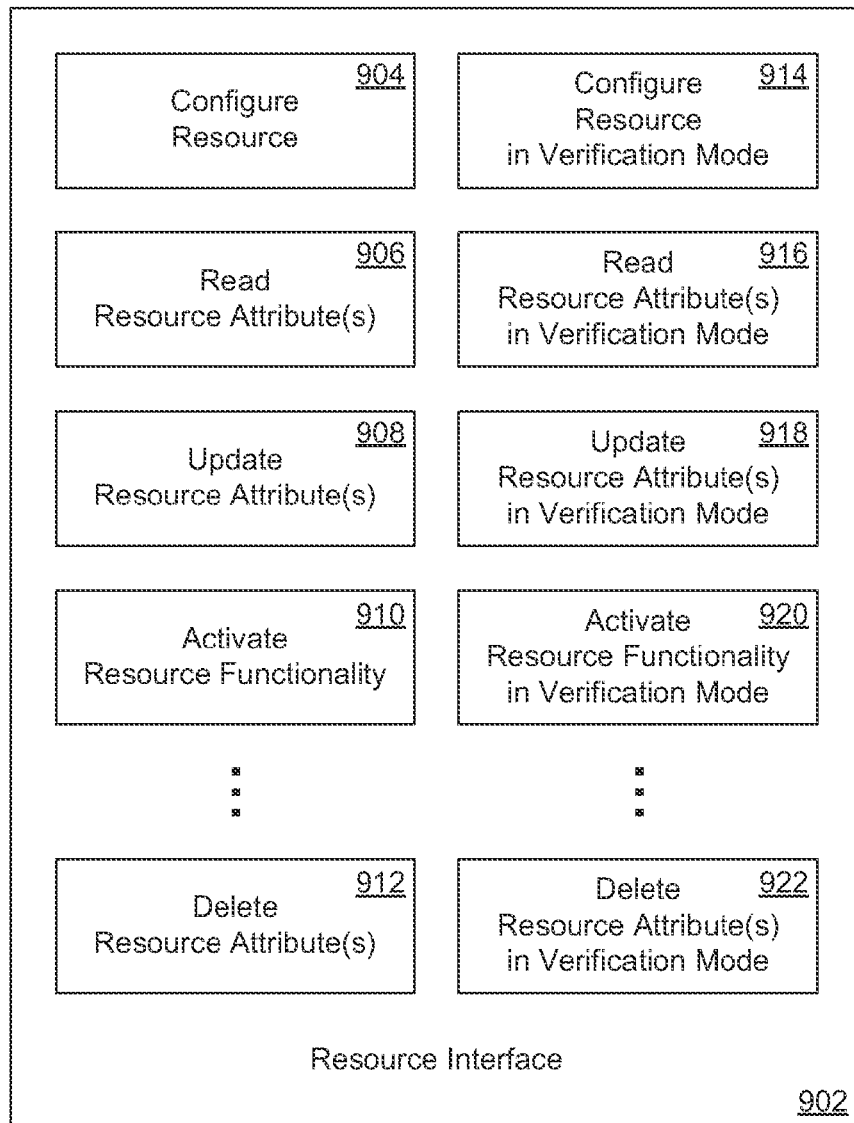
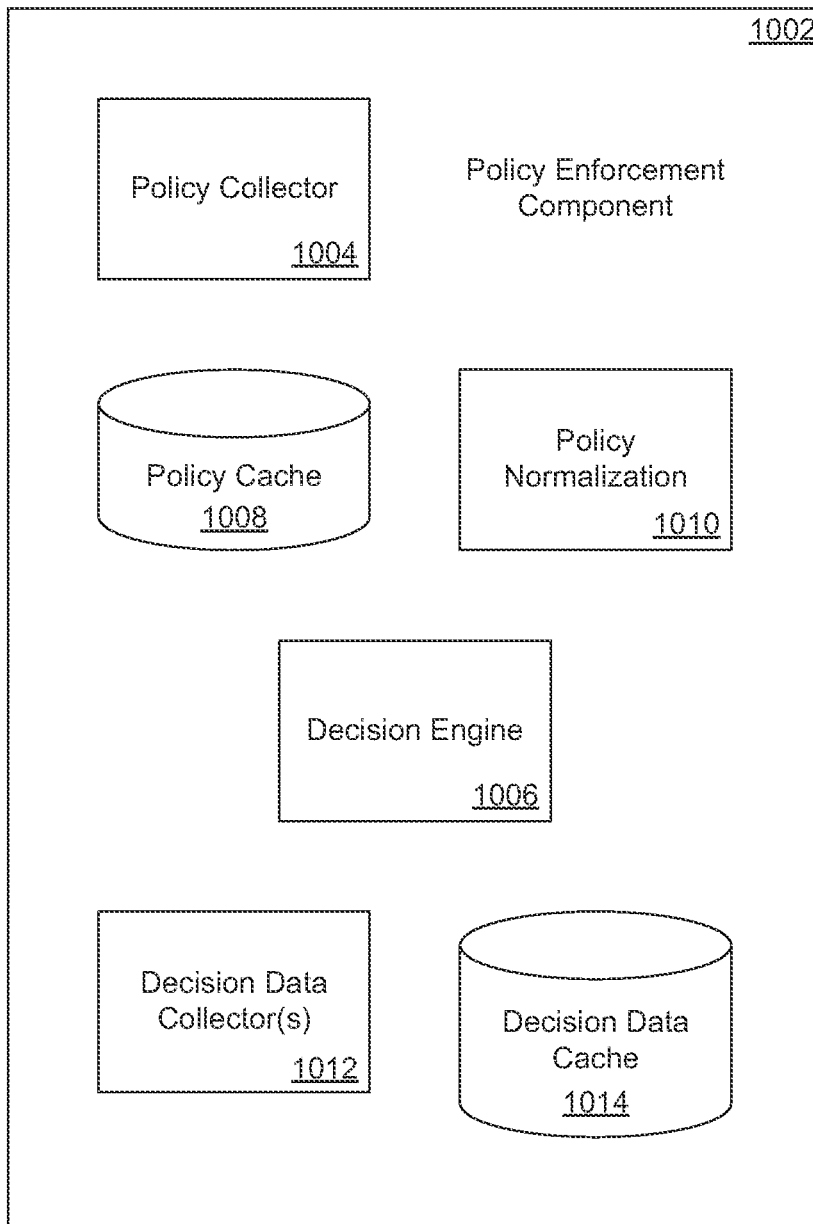
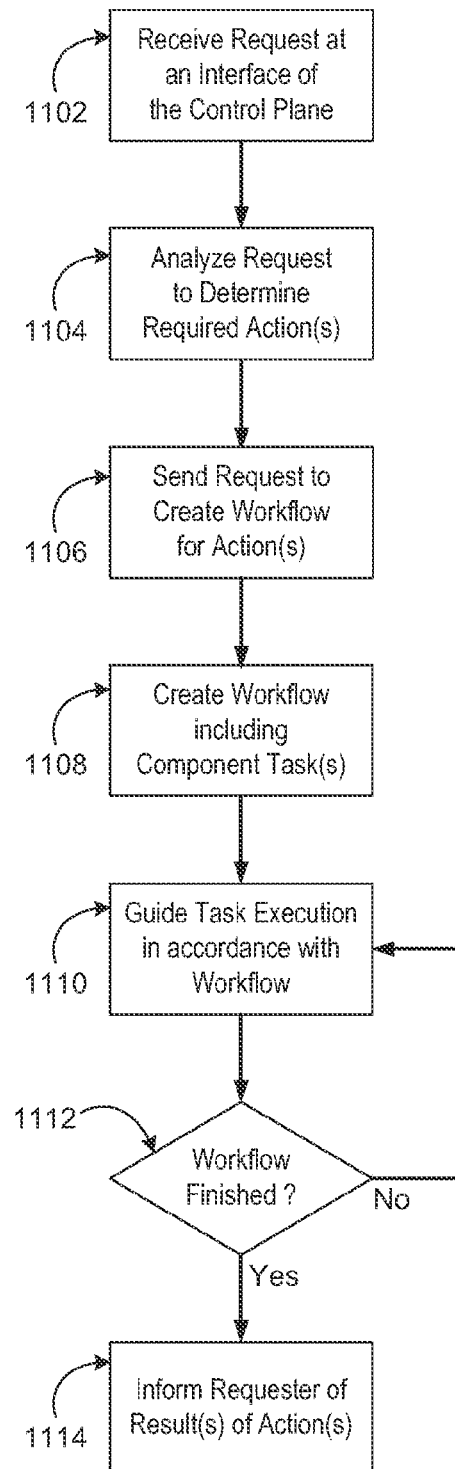


Figure 9

*Figure 10*

*Figure 11*

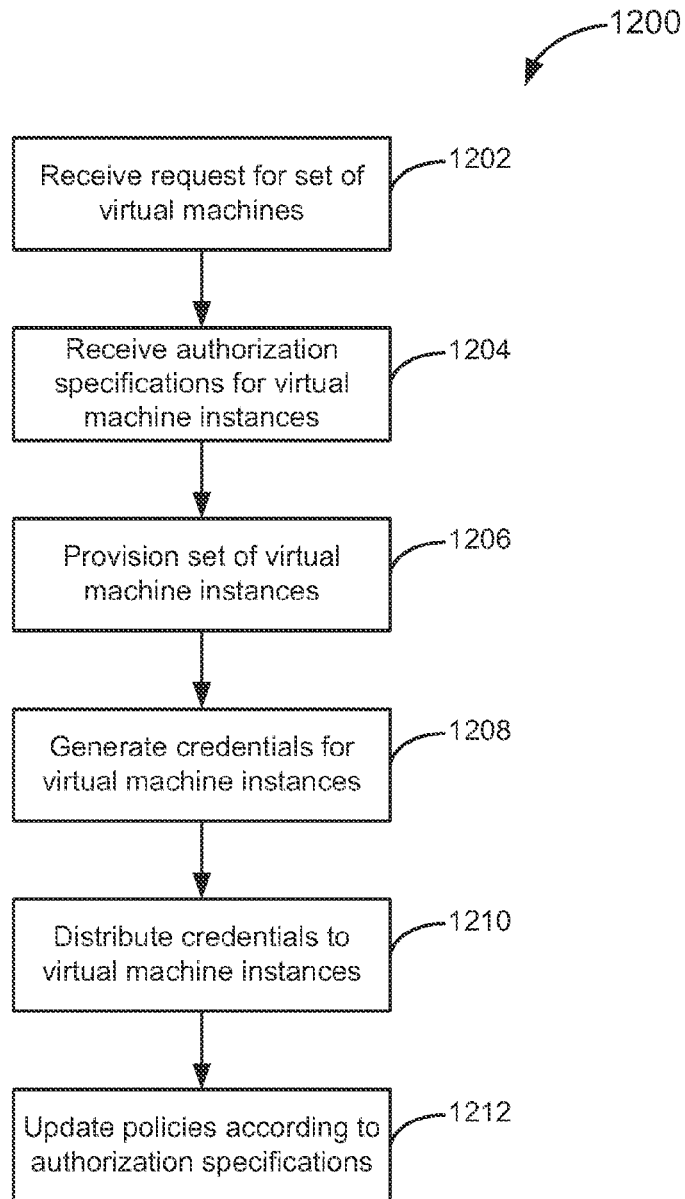


Figure 12

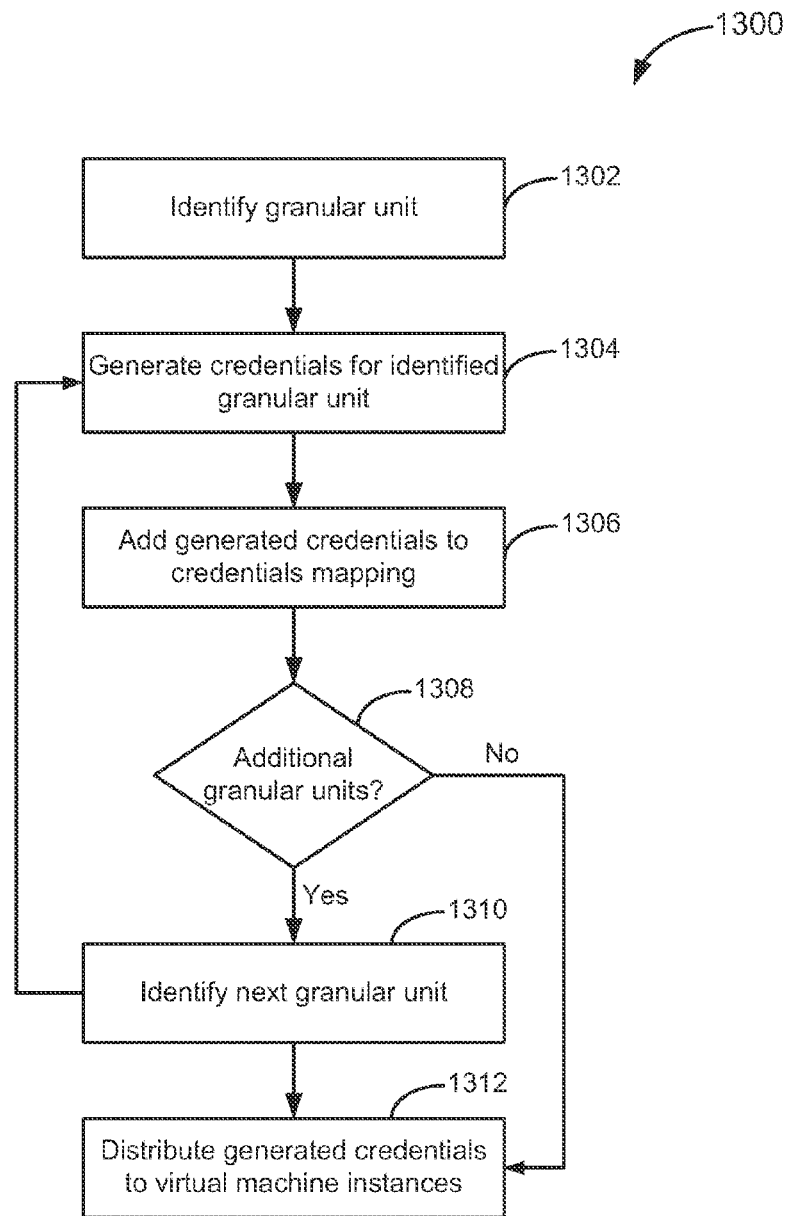


Figure 13

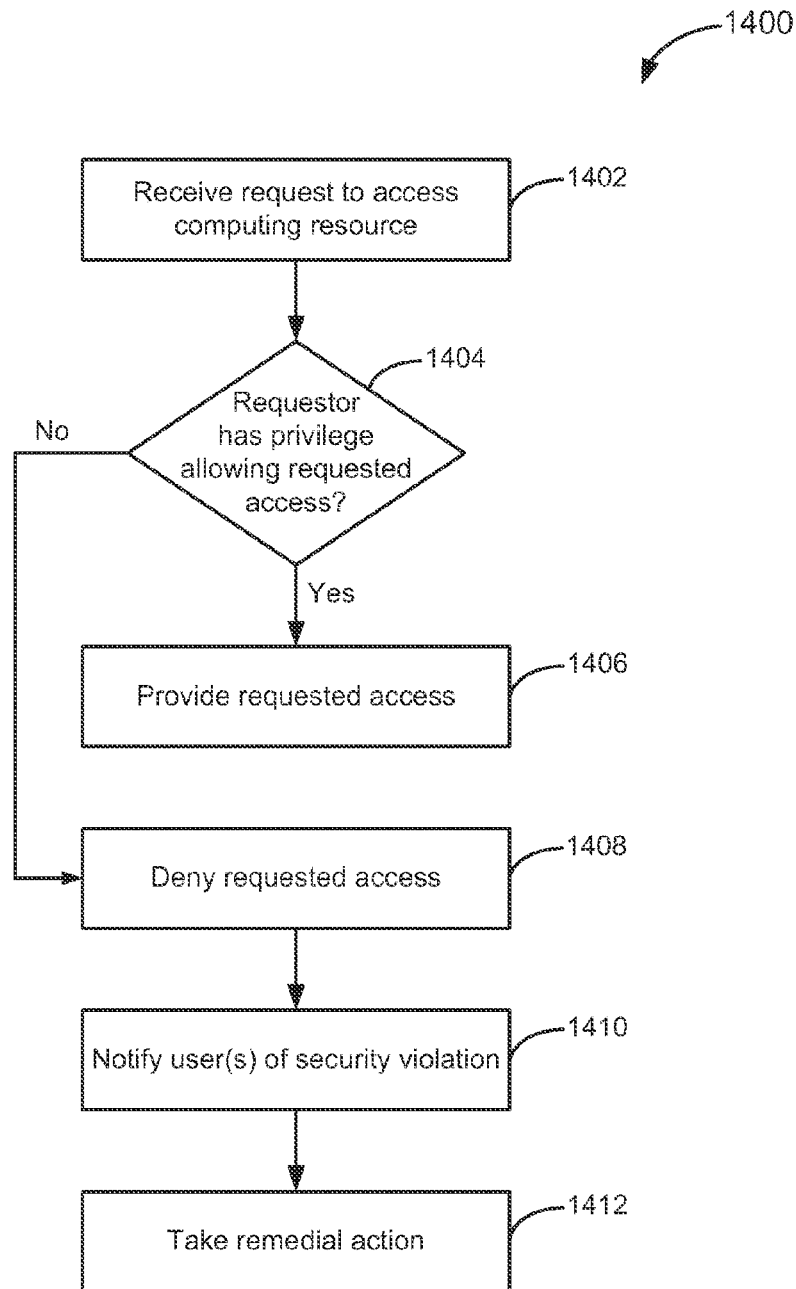


Figure 14

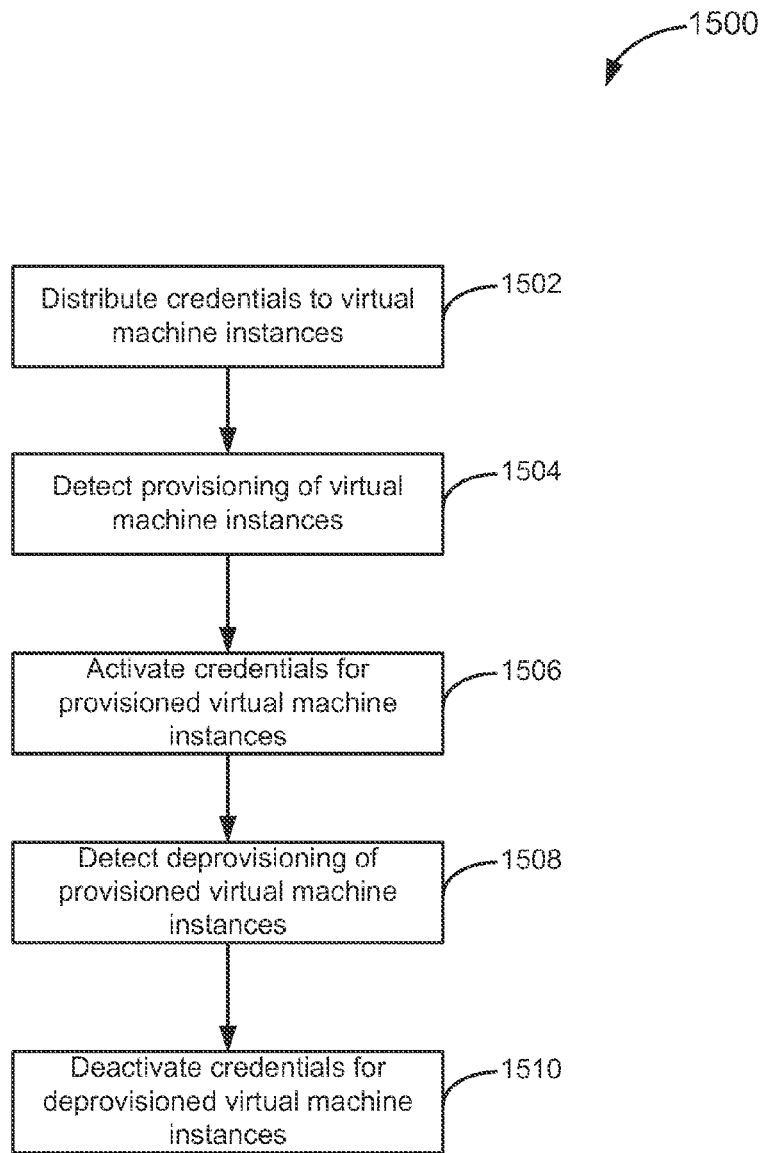


Figure 15

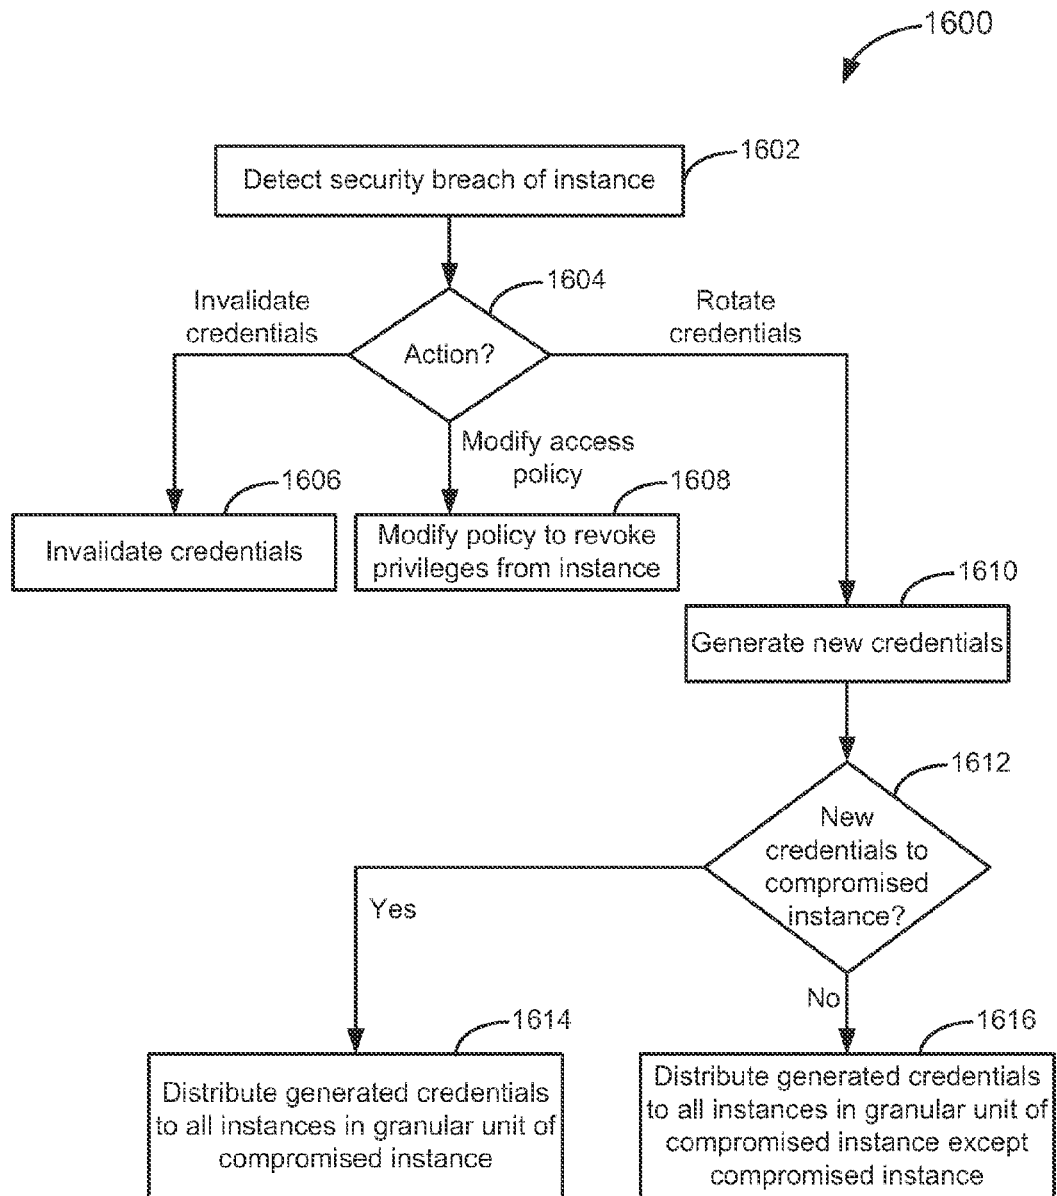


Figure 16

1

TECHNIQUES FOR MANAGING CREDENTIALS IN A DISTRIBUTED COMPUTING ENVIRONMENT

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 12/981,231, filed on Dec. 29, 2010, entitled “TECHNIQUES FOR CREDENTIAL GENERATION,” the content of which is incorporated by reference herein in its entirety.

BACKGROUND

Modern computing networks provide access to a wide variety of computing resources such as data archives, search engines, data processing, data management, communications, electronic marketplaces, as well as media and entertainment services. As the number and size of such computing resources, and their user communities, have grown and become more sophisticated, a need has arisen to establish increasingly sophisticated usage policies. For example, such policies may include policies that address security, privacy, access, regulatory and cost concerns.

Often, policy enforcement in connection with computing resources involves devices of the computing resources authenticating themselves to access other computing resources. A device may provide credentials to another device that are verified to determine actions that may be taken by the device according to any policies applicable to the device. However, modern computer networks often involve the control and administration of multiple devices, including multiple virtual devices. In addition, devices in computing networks are often subject to various attempts at unauthorized access. If an attempt at unauthorized access is successful, there may exist a possibility that credentials used by a compromised device could be used for additional unauthorized access, which potentially could have harmful effects. As a result of the complexities of modern computing networks and the risks to which devices in the networks are subjected, effective management of credentials can be a difficult endeavor.

BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

FIG. 1 is a schematic diagram illustrating an example environment for implementing aspects in accordance with at least one embodiment;

FIG. 2 is a schematic diagram depicting aspects of an example virtual computing resources architecture in accordance with at least one embodiment;

FIG. 3 is a schematic diagram depicting aspects of an example virtual resource provider control plane in accordance with at least one embodiment;

FIG. 4 is a schematic diagram depicting aspects of an example normal form policy in accordance with at least one embodiment;

FIG. 5 is a schematic diagram depicting aspects of an example virtual resource provider in accordance with at least one embodiment;

FIG. 6 is a schematic diagram depicting aspects of an example policy management service in accordance with at least one embodiment;

2

FIG. 7 is a schematic diagram depicting aspects of an example verification mode service in accordance with at least one embodiment;

FIG. 8 is a schematic diagram depicting aspects of an example virtual resource service in accordance with at least one embodiment;

FIG. 9 is a schematic diagram depicting aspects of an example resource interface in accordance with at least one embodiment;

FIG. 10 is a schematic diagram depicting aspects of an example policy enforcement component in accordance with at least one embodiment;

FIG. 11 is a flowchart depicting example steps for workflow management in accordance with at least one embodiment;

FIG. 12 is a flowchart depicting example steps for credential management in accordance with at least one embodiment;

FIG. 13 is a flowchart depicting example steps for generating credentials, which may be used in connection with the process illustrated in FIG. 12, in accordance with at least one embodiment;

FIG. 14 is a flowchart depicting example steps for using credentials in accordance with at least one embodiment;

FIG. 15 is a flowchart depicting example steps for managing credential states in accordance with at least one embodiment; and

FIG. 16 is a flowchart depicting example steps for taking remedial action in connection with a security breach in accordance with at least one embodiment.

Same numbers are used throughout the disclosure and FIGS. to reference like components and features, but such repetition of numbers is for purposes of simplicity of explanation and understanding, and should not be viewed as a limitation on the various embodiments.

DETAILED DESCRIPTION

In the following description, various embodiments will be described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

Techniques described and suggested herein include systems and methods for managing credentials. The credentials may be used to allow computing resources requesting access to other computing resources of a computing network to access such resources. In particular, the credentials may be used as information useful in verifying that a computing resource requesting access to another computing resource is allowed the requested access. As an example, one or more virtual machine instances may provide credentials to access another computing resource, such as a data storage computing resource and/or to control, manage, and/or otherwise operate other virtual machine instances. Credentials may be provided in connection with every request to access computing resources or in other ways, such as to establish a connection with a resource to allow access for a period of time. In an embodiment, various resources are accessible through application programming interfaces (APIs) and credentials may be used by one or more virtual machine instances in connection with making the API calls to access other computing resources. One or more policies may specify actions, such as particular API calls which are

permitted or forbidden. One or more policies may condition one or more actions on possession of particular credentials.

In an embodiment, credentials are generated and distributed to computing resources of a security domain. The security domain may include a collection of computing resources that are collectively administrated, configured, managed, and/or operated. The security domain may include, for example, computing resources of a remote computing services provider that are operated on behalf of and that are managed by a customer who remotely provides instructions for the resources' configuration and operation. Thus, in this example, the credentials may be generated and distributed to a plurality of virtual machines operated on behalf of the customer. Generally, the credentials may be generated and distributed to any set of collectively managed computing resources. The credentials themselves may be any information that may be used to access computing resources. An example of credentials includes a unique identifier string for a set of one or more virtual machine instances. Other types of credentials may also be used, such as keys for one or more public key cryptography authentication and/or other processes. Credentials may also encode metadata about associated computing resources. For instance, using the example of a virtual machine instance, credentials may encode an identifier of an owner of the instance, software installed on the instance, a machine image used to instantiate the instance, an operating system of the instance, one or more software licenses attached to the instance's machine image, an Internet protocol (IP) or other identifier of the instance, and/or other information.

In accordance with an embodiment, different credentials are generated for each of a plurality of subsets of a set of computing resources. The subsets may be defined in various ways and may be defined and/or selected by a user of the computing resources. One or more subsets may be defined, for example, based at least in part on a class of host used to instantiate one or more virtual machine instances. As another example, one or more subsets may be defined based at least in part on members of the subsets being part of logical groupings and/or clusters, on geographic location, on user-assigned tags, and/or other characteristics. Generally, the subsets may be defined in any suitable manner including, but not limited to, explicit user assignment of the subsets. In some embodiments, one or more of the subsets may each consist of a single virtual machine instance. Generation of the credentials may be made according to granularity information received from a user. The granularity information may be determinative of the subsets, such as by specifying information by which the subsets are defined. The granularity information may be received in connection with a user request to provision virtual computing resources.

Other information may also be received from a user in addition to granularity information. As an example, information stating and/or defining one or more policies may be received from a user. The policies may be used to determine actions and the conditions under which the actions may be taken. In an embodiment, the policies are used to determine whether one or more virtual machine instances having credentials have privileges to perform one or more actions. In some instances, possession of particular credentials and/or a class of credentials may be part of a set of one or more conditions for policies. In an embodiment, computing resources having credentials distributed in accordance with the various embodiments described herein submit requests to perform one or more actions in connection with one or more other computing resources.

The various embodiments described and suggested herein provide for effective management of credentials used by various computing resources. For instance, if a security breach is detected in connection with one or more virtual computing resources, credentials used by the one or more virtual computing resources may be invalidated to prevent future security breaches using the credentials. Therefore, if unauthorized access to a virtual machine instance occurs, credentials used by the virtual machine instance (and perhaps other virtual machine instances) may be prevented from being used in the future. Credentials shared by one or more virtual computing instances may also be invalidated in other ways, including, but not limited to, according to a schedule for rotating credentials and/or receipt of user-generated instructions for invalidating credentials. In addition, credentials used by one or more virtual machine resources may be at least temporarily inactivated in order to increase security. Inactivation of credentials may be based at least in part on a change of state of a virtual computing resource. As an example, if operation of one or more virtual machine instances sharing credentials is suspended, the credentials may be inactivated for a time period corresponding to the suspension of the one or more virtual machine instances. Other changes of state may also cause inactivation of credentials.

As noted above, various embodiments of the present disclosure make use of policies. Policies may be specified with a convenient policy specification language and/or user interface. However, such user-specified policies may be inefficient for the purposes of policy enforcement, particularly as the set of user-specified policies grows large. In at least one embodiment, user-specified policies may be transformed and/or processed into a normal policy form ("normal form"). The normal form may include a defined set of policy elements, each with a defined set of possible values. For example, the set of policy elements may include an "actions" element corresponding to a set of actions governed by the policy, a "principals" element corresponding to a set of principals or actors with respect to the set of actions, a "resources" element corresponding to a set of resources or subjects of the action, and/or a "conditions" element corresponding to a set of conditions to be satisfied before the policy is enforced.

The set of normal form policies may be efficiently optimized, for example, optimized for efficient indexing and/or optimized with respect to set size. Normal form policies may be split and/or joined with respect to policy elements, and duplicates efficiently removed. Normal form policies may be efficiently indexed with respect to policy element values and/or particular subsets of policy element values. For example, when the set of normal form policies are to be enforced with respect to a variety of computing resources, the set may be indexed by computing resource, computing resource type, computing resource server, and suitable combinations thereof. Subsets of the set of normal form policies may be distributed to multiple policy enforcement components based on the indexing and/or based on a relevance of particular policy subsets to particular policy enforcement components.

A virtual resource provider capable of provisioning a variety of virtual computing resources may include a centralized policy management service that maintains the set of normalized policies and distributes appropriate subsets to multiple policy enforcement components distributed throughout the virtual resource provider. Alternatively, or in addition, normal form policies may be maintained local to particular policy enforcement components and/or associated

5

virtual computing resource implementation services. As a further alternative, policies may be provided along with action requests. In each case, a policy may be compatible with the normal form, and be associated with and/or accompanied by cryptographic credentials that authenticate and/or authorize the policy and/or an associated user. For example, the user may be authorized to set policies with respect to a particular set of computing resources. Such policies themselves are an example of a computing resource that may be subject to policy control.

Requests with respect to resources may be submitted to one or more policy enforcement components. For example, a request may be submitted to one or more policy enforcement components locally maintaining one or more policies that are relevant to the request. A distributed policy enforcement system with multiple policy enforcement components may have a “sandbox” mode and/or verification mode for processing requests. An authorized user may activate the verification mode with an indicator incorporated into a service interface, with an explicit verification mode service interface, and/or with a cryptographic verification mode token, the acquisition of which is subject to policy control. During the verification mode, requests may be evaluated with respect to policies as usual while requested actions are inhibited.

Such evaluations may depend on a set of decision data including user identity, a user’s group memberships, protocol level details such as network source address of the request and protocol utilized to transmit the request, geographic regions associated with the request such as source geographic location of the request, regulatory classifications associated with the request, date and time. Authorized users may substitute pieces of decision data with specified values, for example, to enable an administrator to verify that a particular set of policies will govern the actions of a particular user or group of users as intended, and/or to diagnose problems. In verification mode, evaluation results, relevant policies, and decision data utilized during request evaluation may be collected, filtered and reported at a variety of levels of detail. The reported set of relevant policies may include normal form policies utilized to evaluate a request and/or a corresponding set of user-specified policies. Originating user-specified policies may be tracked during the policy normalization process to enable reporting of user-specified policies in verification mode. The reported set of decision data may include decision data not otherwise accessible to the report recipient.

Various approaches may be implemented in various environments for various applications. For example, FIG. 1 illustrates aspects of an example environment 100 for implementing aspects in accordance with various embodiments. As will be appreciated, although a Web-based environment may be utilized for purposes of explanation, different environments may be utilized, as appropriate, to implement various embodiments. The environment 100 shown includes both a testing or a development portion (or side) and a production portion. The production portion includes an electronic client device 102, which may include any appropriate device operable to send and receive requests, messages, or information over an appropriate network 104 and convey information back to a user of the device 102. Examples of such client devices include personal computers, cell phones, smart phones, handheld messaging devices, laptop computers, tablet computers, set-top boxes, personal data assistants, electronic book readers, and the like.

The network 104 may include any appropriate network, including an intranet, the Internet, a cellular network, a local

6

area network, a wide area network, a wireless data network, or any other such network or combination thereof. Components utilized for such a system may depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network may be enabled by wired or wireless connections, and combinations thereof. In this example, the network 104 includes the Internet, as the environment includes a Web server 106 for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be utilized as would be apparent to one of ordinary skill in the art.

The illustrative environment 100 includes at least one application server 108 and a data store 110. It should be understood that there may be several application servers, layers, or other elements, processes, or components, which may be chained or otherwise configured, which may interact to perform tasks such as obtaining data from an appropriate data store. As used herein the term “data store” refers to any device or combination of devices capable of storing, accessing, and/or retrieving data, which may include any combination and number of data servers, databases, data storage devices, and data storage media, in any standard, distributed, or clustered environment.

The application server 108 may include any appropriate hardware and software for integrating with the data store as needed to execute aspects of one or more applications for the client device 102, and may even handle a majority of the data access and business logic for an application. The application server 108 provides access control services in cooperation with the data store 110, and is able to generate content such as text, graphics, audio, and/or video to be transferred to the user, which may be served to the user by the Web server 106 in the form of HTML, XML, or another appropriate structured language in this example.

The handling of all requests and responses, as well as the delivery of content between the client device 102 and the application server 108, may be handled by the Web server 106. It should be understood that the Web and application servers 106, 108 are not required and are merely example components, as structured code discussed herein may be executed on any appropriate device or host machine as discussed elsewhere herein. Further, the environment 100 may be architected in such a way that a test automation framework may be provided as a service to which a user or application may subscribe. A test automation framework may be provided as an implementation of any of the various testing patterns discussed herein, although various other implementations may be utilized as well, as discussed or suggested herein.

The environment 100 may also include a development and/or testing side, which includes a user device 118 allowing a user such as a developer, data administrator, or tester to access the system. The user device 118 may be any appropriate device or machine, such as is described above with respect to the client device 102. The environment 100 may also include a development server 120, which functions similar to the application server 108 but typically runs code during development and testing before the code is deployed and executed on the production side and becomes accessible to outside users, for example. In some embodiments, an application server may function as a development server, and separate production and testing storage may not be utilized.

The data store **110** may include several separate data tables, databases, or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store **110** illustrated includes mechanisms for storing production data **112** and user information **116**, which may be utilized to serve content for the production side. The data store **110** also is shown to include a mechanism for storing testing data **114**, which may be utilized with the user information for the testing side. It should be understood that there may be many other aspects that are stored in the data store **110**, such as for page image information and access right information, which may be stored in any of the above-listed mechanisms as appropriate or in additional mechanisms in the data store **110**.

The data store **110** is operable, through logic associated therewith, to receive instructions from the application server **108** or development server **120**, and obtain, update, or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store **110** might access the user information **116** to verify the identity of the user, and may access the catalog detail information to obtain information about items of that type. The information then may be returned to the user, such as in a results listing on a Web page that the user is able to view via a browser on the user device **102**. Information for a particular item of interest may be viewed in a dedicated page or window of the browser.

Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server, and typically will include a computer-readable medium storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available, and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

The environment **100** in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in FIG. 1. Thus, the depiction of the system **100** in FIG. 1 should be taken as being illustrative in nature, and not limiting to the scope of the disclosure.

In at least one embodiment, one or more aspects of the environment **100** may incorporate and/or be incorporated into a virtual computing resources architecture. FIG. 2 depicts aspects of an example virtual computing resources architecture **200** in accordance with at least one embodiment. The example virtual computing resources architecture **200** includes a virtual resource provider **202** enabling various clients **204-206** to interact with provisioned computing resources **216** over a network **214**. The provisioned computing resources **216** may include multiple types of virtual resources **218-220** such as virtual computing systems and clusters, virtual file system volumes, virtual private networks, data object stores, notification services, and suitable combinations thereof. The ellipsis between the types of virtual resources **218** and **220** indicates that the virtual resource provider **202** may maintain any suitable number of such computing resource types and/or instances. Ellipses are used similarly throughout the drawings.

The virtual resource provider **202** may be implemented, at least in part, with server computers such as the Web server **106** and the application server **108** described above with reference to FIG. 1, and one or more data stores such as the data store **110** of FIG. 1, interconnected by a relatively high speed data network (not shown in FIG. 2). The server computers and/or data store(s) implementing the virtual resource provider **202** may include different types and/or configurations of server computers and/or data stores allocated to different virtual resource types **218-220** and to a control plane **222** of the virtual resource provider **202**. The control plane **222** may include multiple user interfaces **224-226** that enable the clients **204-206** to interact with the virtual resource provider **202**, including provisioning and interacting with the virtual resources **218-220**, as well as setting policies with respect to the virtual resources **218-220**. FIG. 3 depicts further aspects of a control plane **302** in accordance with at least one embodiment.

The user interfaces **304** of the control plane **302** may include any suitable type of user interface. For example, the user interface **304** may include graphical user interfaces (GUIs), Web-based interfaces, programmatic interfaces such as application programming interfaces (APIs) and/or sets of remote procedure calls (RPCs) corresponding to interface elements, messaging interfaces in which the interface elements correspond to messages of a communication protocol, and/or suitable combinations thereof. The user interfaces **304** may include multiple resource interfaces **306-308** enabling user interaction with the virtual resources **218-220** (FIG. 2). For example, there may be one of the resource interfaces **306-308** for each of the types of virtual resources **218-220**. As another example, one or more of the resource interfaces **306-308** may enable interaction with multiple virtual resource types **218-220**. Alternatively, or in addition, one or more of the types of virtual resources **218-220** may incorporate interfaces enabling direct interaction (e.g., by the clients **204-206**) without mediation by one of the user interfaces **304**.

The user interfaces **304** may further include a provisioning interface **310** enabling authorized users to provision, configure, re-configure and/or de-provision (collectively, "provision") the virtual resources **218-220** (FIG. 2). Although the example control plane **302** includes one such provisioning interface **310**, embodiments may include any suitable number of such provisioning interfaces, including a provisioning interface for each type of virtual resource **218-220**. Alternatively, or in addition, the resource interface **306** corresponding to a particular type of virtual resource **218** may incorporate the provisioning interface for that virtual resource type **218**. The user interfaces **304** of the control plane **302** may still further include a policy management interface **312** enabling establishment, viewing, editing and/or deletion (collectively, "maintenance") of virtual resource provider **202** policies including policies with respect to the virtual resources **218-220** and with respect to policy management. Again, although the example control plane **302** includes one such policy management interface **312**, embodiments may include any suitable number of policy management interfaces including a policy management interface for each type of virtual resource **218-220**. The resource interface **306** corresponding to a particular type of virtual resource **218** may incorporate the policy management interface for that virtual resource type **218**.

The control plane **302** may include a policy management component **314** configured at least to provide centralized policy management for the virtual resource provider **202** (FIG. 2). The policy management component **314** may

receive user-specified policies, for example, specified with the policy management interface **312**, and transform user-specified policies into a normal form. The policy management component **314** may then optimize the set of normal form policies, for example, for indexing and/or with respect to set size, and index the set of normal form policies based at least in part on one or more policy elements of the normal form. The policy management component **314** may then distribute the optimized normal form policies to a set of policy enforcement components **316-318** in a manner based at least in part on the index.

The control plane **302** may further include a verification mode component **320** configured at least to facilitate verification mode reporting. The verification mode component **320** may collect request evaluation results, normal form policies and/or decision data, map normal form policies to user-specified policies, and generate reports at specified levels of detail. For example, verification mode report detail levels may include a low detail level corresponding to an evaluation result and basic contextual information, a medium detail level that adds a set of policies considered to determine the evaluation result, a high detail level that further adds a set of decision data considered to determine the evaluation result, and a very high detail level that further adds virtual resource provider **202** (FIG. 2) operational information that may assist anomalous behavior diagnosis (e.g., debugging). For example, such operational information may include an identification of the set of policy enforcement components **316-318** involved in evaluating a request associated with the report. In at least one embodiment, verification mode may be a mechanism by which authorized users are permitted access to such virtual resource provider **202** operation information. Verification mode may correspond to a sandbox mode in the software development and/or testing sense, enabling configuration verification and/or testing without possible high cost consequences (e.g., in terms of financial and/or resource costs).

The verification mode component **320** may be further configured to process requests for verification mode tokens (e.g., cryptographic tokens), and to authenticate such tokens. For example, a verification mode token may be submitted with a request in order to indicate that the request is to be processed in verification mode. Alternatively, or in addition, verification mode tokens may authorize substitution of decision data (or at least a portion of the decision data) utilized when evaluating a request with respect to a set of policies. For example, the verification mode token may authorize substitution of decision data related to user identity, thus enabling an administrator to verify that a particular policy set has an intended effect with respect to a particular user.

The control plane **302** may include a workflow management component **322** configured at least to establish and maintain workflows such as resource workflows, provisioning workflows and/or policy management workflows established responsive to user interaction with the resource interfaces **306-308**, the provisioning interface **310**, and/or the policy management interface **312**, respectively. Workflows may include one or more sequences of tasks to be executed to perform a job, such as virtual resource configuration, provisioning or policy management. A workflow, as the term is used herein, is not the tasks themselves, but a task control structure that may control flow of information to and from tasks, as well as the order of execution of the tasks it controls. For example, a workflow may be considered a state machine that can manage and return the state of a process at any time during execution. Workflows may be created from workflow templates. For example, a policy management

workflow may be created from a policy management workflow template configured with parameters by the policy management component **314**.

The workflow management component **322** may modify, further specify and/or further configure established workflows. For example, the workflow management component **322** may select particular implementation resources of the virtual resource provider **202** to execute and/or be assigned to particular tasks. Such selection may be based at least in part on the computing resource needs of the particular task as assessed by the workflow management component **322**. As another example, the workflow management component **322** may add additional and/or duplicate tasks to an established workflow and/or reconfigure information flow between tasks in the established workflow. Such modification of established workflows may be based at least in part on an execution efficiency analysis by the workflow management component **322**. For example, some tasks may be efficiently performed in parallel, while other tasks depend on the successful completion of previous tasks.

The normal form of a policy may include a defined set of policy elements. FIG. 4 depicts aspects of an example normal form policy **402** in accordance with at least one embodiment. The policy **402** includes a policy identifier **404** uniquely identifying the policy **402**, one or more effects **406** of the policy **402**, reference to one or more principals **408** or actors governed by the policy, a specification of one or more actions **410** governed by the policy **402**, reference to one or more resources **412** governed by the policy **402**, a specification of one or more conditions **414** to be satisfied before the policy **402** is enforced, and a set of element modifiers **416**, **418**, **420**, **422** corresponding to the policy elements **408**, **410**, **412**, **414**, respectively. The policy **402** may have an alphanumeric and/or Unicode representation and/or encoding sometimes called the policy document.

The policy identifier **404** may be any suitable identifier uniquely identifying the policy **402**. For example, the policy identifier **404** may correspond to an alphanumeric and/or Unicode (collectively, "alphanumeric") string or integer number. The policy identifier **404** may be utilized to reference the policy **402**. Examples of policy effects include permit and deny. The effect(s) **406** element of the policy **402** may specify such policy effects. For example, a particular policy may permit one or more principals **408** to take one or more actions **410** with respect to one or more resources **412**, while another policy may deny a set of actions **410** to a set of principals **408**.

The principal(s) **408** element of the policy **402** may specify one or more entities known to the virtual resource provider **202** (FIG. 2) that are capable of making requests of the virtual resource provider **202**. Such entities may include users having a user account with the virtual resource provider **202**, customers having a commercial account (e.g., a cost-tracking account) with the virtual resource provider **202**, and groups of users and/or customers including role-based groups such as administrators. Virtual machine instances or other virtual resources, other resources, and/or groups thereof may be also be principals. Such entities may be specified with any suitable identifier including user identifiers, customer account numbers, group identifiers, and alphanumeric strings.

The action(s) **410** element of the policy **402** may specify one or more actions capable of being performed by the virtual resource provider **202** (FIG. 2). Such actions may include actions capable of being performed by the virtual resources **218-220** of the virtual resource provider **202** and/or actions that may be requested with the user interfaces

224-226 of the virtual resource provider 202, including policy management actions. In at least one embodiment, actions specified by the action(s) 410 element correspond to elements of the user interfaces 224-226. Such actions may be specified with any suitable action identifier including interface element identifiers and alphanumeric strings.

The resource(s) 412 element of the policy 402 may specify one or more resources of the virtual resource provider 202 (FIG. 2) for which one or more of the action(s) 410 is valid. Such resources may include virtual resources 218-220, implementation resources, policies, and verification mode tokens. Such resources may be specified with any suitable resource identifier including resource identifiers in accordance with a uniform resource identifier (URI) standard such as Berners-Lee et al., "Uniform Resource Identifier (URI): Generic Syntax," Request for Comments (RFC) 3986, Network Working Group, January 2005.

The condition(s) 414 element of the policy 402 may specify a set of conditions to be satisfied before the policy 402 is enforced. For example, the condition(s) 414 may specify the conditions under which the principal(s) 408 are permitted to perform the action(s) 410 with respect to the resource(s) 412. Such conditions may be specified with any suitable condition specification language including suitable programming languages, and may include compound conditions, for example, specified with Boolean operators. Condition parameters may include any suitable data available to the virtual resource provider 202 (FIG. 2). Condition parameter examples include environmental data such as calendar date and time of day, and request-associated data such as originating network address, originating geographical location, originating political and/or administrative division and communication protocol employed.

By default, the policy 402 effect(s) 406 may be enforced when the specified principal(s) 408 request the specified action(s) 410 with respect to the specified resource(s) 412 and the specified set of condition(s) 414 are satisfied. However, the element modifiers 416-422 may indicate that the corresponding policy elements 408-414 specify exceptions. That is, that the policy 402 effect(s) 406 be enforced unless the specified principal(s) 408 are the actors, unless the specified action(s) 410 are requested, unless the specified resource(s) 412 are the subject of action and/or unless the specified set of conditions is satisfied. In addition, policy element 408-414 specifications may include "wildcard" values. For example, a "*" wildcard value may match all and/or any valid values of the policy element 408-414. Alternatively, or in addition, wildcard values may be matched contextually. For example, a "*" wildcard value may match all and/or any values of the policy element 408-414 that are valid with respect to a particular user, customer, group, and/or other context of a particular request.

As described above with reference to FIG. 3, the policy management component 314 may transform user-specified policies into normal form policies, and optimize sets of normal form policies. The optimized set of normal form policies may differ significantly from the supplied set of user-specified policies. A single user-specified policy can result in the addition of multiple normal form policies. A single normal form policy can participate in enforcing multiple user-specified policies. In at least one embodiment, each normal form policy 402 maintains a set of references to "parent" policies 424 which the normal form policy 402 participates in enforcing. For example, the set of parent policies 424 may be updated during user-specified policy transformation and normal form policy set optimization operations.

The virtual resource provider 202 may be implemented as a collection of networked services. FIG. 5 depicts aspects of an example virtual resource provider 502 implemented in accordance with at least one embodiment. The virtual resource provider 502 includes multiple virtual resource services 504-506, a policy management service 508, an authentication service 510, a set of decision data providers 512, and a verification mode service 514, each interconnected by a network 516. The arrow 518 indicates that the network 516 is connected to the network 214 of FIG. 2.

Each virtual resource service 504, 506 may maintain a set of provisioned resources 520-522, 524-526 and incorporate a resource interface 528, 530. For example, each virtual resource service 504-506 may maintain one type of virtual resource 218-220 as described above with reference to FIG. 2, and incorporate a corresponding one of the resource interfaces 306-308 described above with reference to FIG. 3. The policy management service 508 may incorporate a policy management interface 532 corresponding to the policy management interface 312 of FIG. 3.

The policy management service 508 may act as a centralized policy management service for the virtual resource provider 502, managing, transforming, optimizing and distributing one or more sets of policies 534 to other services 504-506, 510-514 of the virtual resource provider 502. In addition, in at least one embodiment, services 504-506, 510-514 of the virtual resource provider 502 other than the policy management service 508 may maintain policies. For example, the virtual resource services 504-506 may each maintain a set of policies 536-538, 540-542 that are relevant to the respective set of provisioned resources 520-522, 524-526. Such policies 536-538, 540-542 may be established, viewed, edited and deleted with policy management interface 312 (FIG. 3) functionality integrated into the associated resource interface 528-530. Such policies 536-538, 540-542 may be compatible with the policies 534 maintained by the policy management service 508. For example, the policies 536-538, 540-542 may each have a form corresponding to the normal form policy 402 (FIG. 4). Further in addition, in at least one embodiment, policies may be maintained at a location external to the virtual resource provider 502, and supplied for enforcement as part of interactions with the services 504-514 of the virtual resource provider 502. For example, request 544 incorporates such a policy 546.

The authentication service 510 may provide authentication services for the virtual resource provider 502. For example, the authentication service 510 may authenticate an integrity and/or an authenticity of policies 534-542, 546. In addition, the authentication service 510 may authenticate an identity of clients 204-206 (FIG. 2), customers and/or users of the virtual resource provider 502. For example, the policy management interface 532 and/or the resource interfaces 528-530 may utilize the authentication service 510 to establish and/or verify an identity and/or authority of a user with respect to policies 534-542, 546 and/or provisioned resources 520-526. The authentication service 510 may utilize any suitable authentication mechanism, including any suitable cryptographic authentication mechanism, to provide the authentication services. Such authentication mechanisms are well known in the art and need not be detailed here.

The decision data providers 512 may provide data required by policy enforcement components 316-318 (FIG. 3) to evaluate requests that are subject to policies. For example, the provisioned resources 520-526 may be associated with various labeling and/or naming schemes ("resource names"), and the decision data providers 512 may

13

include a name resolution provider that maps or resolves such resource names to locations (e.g., network locations) within the virtual resource provider **502**. The decision data providers **512** may further include providers of data required to determine whether the condition(s) **414** (FIG. 4) of the policy **402** are satisfied. For example, a particular policy may deny interaction with a particular resource from a specified set of geographic locations and/or regions (collectively, “geographic locations”). In such a case, the decision data providers **512** may include a geographic resolution provider capable of identifying a set of geographic locations associated with the request **544**.

The verification mode service **514** may provide verification mode functionality (e.g., as described above with reference to the verification mode component **320** of FIG. 3) as a service to the other services **504-512** of the virtual resource provider **502**. For example, the resource interfaces **528-530** and/or the policy management interface **532** may query the verification mode service **514** to authenticate a verification mode token and/or may provide evaluation results for reporting in response to a request. An example verification mode service is described below in more detail with reference to FIG. 7. However, the description first turns to further aspects of the policy management service **508** in accordance with at least one embodiment.

The policy management service **602** of FIG. 6 is an example of the policy management service **508** of FIG. 5, and includes a policy management interface **604** corresponding to the policy management interface **532** of FIG. 5. Users of the virtual resource provider **502** may interact with the policy management interface **604** to specify, view, edit and/or delete policies in a user-specified policy set **606** maintained by the policy management service **602**. The policy management service **602** may further maintain a normalized policy set **608** corresponding to the user-specified set **606**. The normalized policy set **608** may be maintained by a policy normalization component **610**. The policy management service **602** may participate in policy enforcement, for example, the policy management service **602** may include a policy enforcement component **616** corresponding to one of the policy enforcement components **316-318** of FIG. 3. In particular, the policy enforcement component **616** may enforce policies referencing actions that may be requested with the policy management interface **604**. Furthermore, in at least one embodiment, at least a portion of the normalized policy set **608**, or a copy thereof, is distributed to policy enforcement components located throughout the virtual resource provider **502**. The policy management service **602** may include a policy distribution component **612** configured at least to distribute the normalized policy set **608** throughout the virtual resource provider **502**, for example, based at least in part on a policy index **614**.

Policies in the user-specified policy set **606** may be specified with alphanumeric strings. A first example of such a policy is:

```
permit administrators*
```

corresponding to permission for members of the group with identifier “administrators” to perform all actions with respect to all resources associated with a particular virtual resource provider **502** (FIG. 5) customer account. A second example is:

```
deny jsmith*unless current_time in business_hours
```

corresponding to denying permission for a user with identifier “jsmith” to perform any action outside of business hours. A third example is:

```
permit jsmith to {create, read, write, delete}
data_object_service:zone_1::public/images
```

14

corresponding to permission for the user with identifier “jsmith” to create, read, write and delete data objects at a particular virtual resource having resource identifier “data_object_service:zone_1::public/images”. A fourth example is:

```
permit msmith to configure
```

```
{vcs::example_corp/*, bds::example_corp/*}
```

corresponding to permission for a user with identifier “msmith” to configure all virtual computer systems (provided by a virtual computing service “vcs”) and block data volumes (provided by a block data service “bds”) associated with a customer having customer identifier “example_corp”. Each user-specified policy may be associated with a unique identifier.

The policy normalization component **610** may generate sets of normalized policies corresponding to user-specified policies, and store the normalized policies in the normalized policy set **608**. For example, each policy in the normalized policy set **608** may have a form corresponding to the normal form policy **402** of FIG. 4. The policy normalization component **610** may generate multiple normal form policies corresponding to a user-specified policy, for example, due to decompounding with respect to policy elements **408-414**. The fourth user-specified policy example described above may be decompounded to generate two normal form policies, both with the effect **406** element specifying a permit effect, the principal **408** element specifying the user with identifier “msmith”, and the action **410** element specifying the configure action. For example, the resource **412** element of one of the normal form policies may specify all virtual computer systems associated with the customer identifier “example corp”, the resource **412** element of the other of the normal form policies may specify all block data volumes associated with the customer identifier “example corp”. In this example, the user-specified policy has been decompounded with respect to the resource **412** element. As another example, the policy normalization component **610** may be configured to decompound the first user-specified policy example described above with respect to the action **410** element and/or the resource **412** element.

Wildcards may be first “expanded” into a set of values, for example, based on a context of the user-specified policy, before being decompounded. The decompounding need not be complete with respect to policy element **408-414** (FIG. 4) values. For example, the policy normalization component **610** may be configured to protect particular subsets of policy element **408-414** values (policy element “atoms” or “indivisibles”) from decompounding. As one example, the set of actions corresponding to “{create, read, write, delete}” may be protected from decompounding.

The policy normalization component **610** may reduce a size of the normalized policy set **608**, for example, by reducing redundancy in the normalized policy set **608**. The policy normalization component **610** may detect that a newly added normal form policy is contained within the scope of one or more of the normalized policy set **608**, and delete the newly added normal form policy from the normalized policy set **608**. Normal form policies, such as the normal form policy **402** of FIG. 4, may correspond to sets of points in one or more policy spaces, for example, spaces having dimensions corresponding to the policy elements **408-414** of the normal form, and/or Boolean-valued functions over the policy spaces. Each effect **406** element value may correspond to such a policy space. For example, there may be a permit space and a deny space. The policy enforcement components **316-318** (FIG. 3) may be configured to require an explicit permit in the absence of a deny in

15

order to permit a particular action, that is, to require a point in a permit space without a corresponding point in the deny space. Wildcards may correspond to regions of such policy spaces. Element modifiers **416-422** may be used to define complimentary sets of the policy spaces. The policy normalization component **610** may detect whether the normal form policy **402** is within the scope of a set of normal form policies (e.g., the normalized policy set **608**) by determining whether the set of points corresponding to the normal form policy **402** is contained within the set of points corresponding to the set of normal form policies.

The policy normalization component **610** may further reduce the size of the normalized policy set **608** by compounding and/or recomponding normal form policies in the normalized policy set **608** with respect to policy elements **408-414** (FIG. 4). Some policy element **408-414** values may correspond to multiple other policy element **408-414** values, so that a single normal form policy **402** may correspond to a same policy space region as multiple other normal form policies. For example, a principal **408** element value corresponding to a group may be equivalent to multiple principal **408** element values corresponding to individual users. The policy normalization component **610** may be configured to detect compoundable subsets of the normalized policy set **608** and to compound the detected subsets along policy element **408-414** dimensions to configured granularities (which need not be the coarsest available granularity). Throughout generation and/or transformation of the normalized policy set **608**, the policy normalization component **610** may be configured to maintain the set of parent policies **424** of each normal form policy **402**. For example, the set of parent policies **424** may include the identifier of each user-specified policy that the normal form policy **402** participates in enforcing.

Different regions of the policy space(s) may be enforced by different policy enforcement components **316-318** (FIG. 3). Different policy enforcement components **316-318** may enforce overlapping, and even same, regions of the policy space(s). However, in at least one embodiment, different policy enforcement components **316-318** enforce disjoint regions of the policy space(s). The policy distribution component **612** may determine regions of the policy space(s), and corresponding subsets of the normalized policy set **608**, to assign and/or distribute to the policy enforcement components **316-318**. To facilitate such distribution, the policy distribution component **612** may establish and maintain the policy index **614**. The policy index **614** may index the normalized policy set **608** with respect to policy element **408-414** (FIG. 4) values and/or subsets thereof. For example, the policy index **614** may index the normalized policy set **608** with respect to resource **412** and/or resource **412** type. Policies in the normalized policy set **608** may be multiply indexed, for example, with respect to multiple different policy element **408-414** values and/or value subsets. The policy distribution component **612** may distribute the normalized policy set **608** to the policy enforcement components **316-318** based at least in part on the policy index **614**. For example, sets of indices of the policy index **614** may define policy enforcement domains corresponding to subsets of the normalized policy set **608**, and the policy distribution component **612** may distribute such subsets to the policy enforcement components **316-318** as units. In at least one embodiment, each normal form policy **402** incorporates its indexing information and/or the policy index **614** is incorporated into the normalized policy set **608**.

The distributed nature of policy enforcement in the virtual resource provider **502** (FIG. 5) can complicate efforts to

16

verify that a particular set of user-specified policies has its intended consequences. FIG. 7 depicts aspects of an example verification mode service **702** in accordance with at least one embodiment that can facilitate such verification.

The verification mode service **702** may include a token management component **704** accessible through a token interface **706**. The token management component **704** may issue verification mode tokens in response to authorized requests, record the issued tokens in an issued token database **708**, and authenticate the issued tokens upon request. Following processing of the request **544** in verification mode, a reporting component **710** of the verification mode service **702** may prepare a verification mode report at a specified level of detail. Requests for such reports, as well as level of detail specification, may be made with a reporting interface **712** of the verification mode service **702**. Such report requests may not include all the data required by the reporting component **710**. The verification mode service **702** may further include a data collection component **714** configured to interact with other services **504-512** of the virtual resource provider **502** to obtain the required data.

Policies with respect to actions that may be requested with the token interface **706** and/or the reporting interface may be enforced by a policy enforcement component **716** corresponding to one of the policy enforcement components **316-318** of FIG. 3. For example, such policies may define a set of users of the virtual resource provider **502** (FIG. 5) authorized to be issued verification mode tokens, and may require an authentic verification mode token before providing the verification mode report. The token management component **704** may issue multiple types of tokens, for example, tokens authorizing verification mode reporting may be a separate type from tokens authorizing substitution of decision data during verification mode processing of the request **544**. Tokens authorizing substitution of different subsets and/or types of decision data may also be separate types. Verification mode policies may distinguish between different types of tokens. Generation, issuance, authentication and revocation of cryptographic tokens are well known in the art, and need not be described here in detail.

Requests at the reporting interface **712** may specify a type of report and/or level of reporting detail. For example, the workflow management component **322** (FIG. 3) may generate the request in response to detecting a verification mode request at one of the resource interfaces **306-308**. The reporting component **710** may generate the requested report based on a corresponding report template, utilizing the data collection component **714** to collect data as necessary. Fields that may be included in the verification mode report include one or more results of evaluation of the verification mode request by one or more of the policy enforcement components **316-318**, a subset of the normalized policy set **608** (FIG. 6) and/or the user-specified policy set **606** relevant to the verification mode request and/or utilized during the evaluation by the one or more policy enforcement components **316-318**, decision data utilized during the evaluation, virtual resource provider **502** contextual operating data capable of influencing the evaluation, and/or versions thereof filtered and/or summarized to the specified level of detail. The reporting component **710** may be configured to map a set of normal form policies to a corresponding set of user-specified policies based at least in part on the set of parent policies **424** (FIG. 4) referenced by each normal form policy **402**.

Each virtual resource service **504-506** (FIG. 5) may be implemented with a collection of physical server computers and/or network elements. FIG. 8 depicts aspects of an

17

example virtual resource service **802** in accordance with at least one embodiment. The virtual resource service **802** depicted in FIG. 8 is an example of the virtual resource service **504** of FIG. 5. The example virtual resource service **802** includes multiple virtual resource servers **804-806** and one or more interface servers **808** interconnected by a network **810**. The provisioned resources **520-522** and policies **536-538** of FIG. 5 may be distributed across the virtual resource servers **804-806** of FIG. 8 utilizing any suitable data distribution technique. For example, each virtual resource server **804, 806** may maintain a subset **812-814, 816-818** of the provisioned resources **520-522**, and a corresponding subset **820-822, 824-826** of the policies **536-538**. Such subsets may deliberately overlap for various purposes including data durability and service capability.

The interface server(s) **808** may (collectively) maintain a resource interface **828** and a policy enforcement component **830**. The resource interface **828** may correspond to one of the resource interfaces **306-308** of FIG. 3. Similarly, the policy enforcement component **830** may correspond to one of the policy enforcement components **316-318** of FIG. 3. The policy enforcement component **830** of FIG. 8 may enforce policies referencing actions that may be requested with the resource interface **828** of FIG. 8. Such policies may include policies distributed by the policy management service **602** (FIG. 6), policies **820-826** stored on the virtual resource servers **804-806**, and policies accompanying requests to the resource interface **828** such as the policy **546** of FIG. 5. For example, the resource interface **828** may require a “permit” evaluation by the policy enforcement component **830** before further processing a request.

FIG. 9 depicts aspects of an example resource interface **902** in accordance with at least one embodiment. The resource interface **902** may include any suitable number of interface elements of any suitable type. In this example, the resource interface **902** includes a configure resource interface element **904**, a read resource attribute(s) interface element **906**, an update resource attribute(s) interface element **908**, an activate resource functionality interface element **910**, and a delete resource attribute(s) interface element **912**. In addition, the resource interface **902** includes distinct interface elements **914-922** corresponding to verification mode versions of the interface elements **904-912**. Each interface element **904-922** defines a structured interaction with the provisioned resources **812-818** (FIG. 8) including a request to perform a set of actions with respect to at least one of the provisioned resources **812-818**. In the example resource interface **902**, each interface element **904-922** corresponds to a set of messages of a Web-based services protocol such as extensible markup language (XML) encoded remote procedure calls (e.g., XML-RPC) or a simple object access protocol (e.g., SOAP).

Clients **204-206** (FIG. 2) may utilize the configure resource interface element **904** to request a configuration and/or reconfiguration of one or more of the provisioned resources **812-818** (FIG. 8). For example, if the provisioned resources **812-818** include data object stores, the configure resource interface element **904** may enable clients **204-206** to set a maximum size of a data object store over which they have authority. Clients **204-206** may utilize the read resource attribute(s) interface element **906** to request a read or view of one or more attributes of one or more of the provisioned resources **812-818**. For example, the read resource attribute(s) interface element **906** may enable clients **204-206** to obtain copies of specified data objects from specified data object stores. Clients **204-206** may utilize the update resource attribute(s) interface element **908** to request an

18

update of one or more attributes of one or more of the provisioned resources **812-818**. For example, the update resource attribute(s) interface element **908** may enable clients **204-206** to update specified data objects and/or create specified data objects in specified data object stores.

Some provisioned resources **812-818** (FIG. 8) may have functionality that can be activated, and the activate resource functionality interface element **910** may be utilized to request an activation of that functionality. For example, some types of data object store may have a capability to analyze stored data objects, and the activate resource functionality interface element **910** may enable authorized clients to start, stop, suspend and/or perform the analysis. The delete resource attribute(s) interface element **912** may enable clients **204-206** (FIG. 2) to request a deletion and/or re-initialization of one or more attributes of one or more of the provisioned resources **812-818**. For example, clients **204-206** with sufficient authorization may delete specified data objects from specified data object stores.

In the example resource interface **902**, utilization of the verification mode versions **914-922** of the interface elements **904-912** indicates that the corresponding request is to be processed in the verification mode (is a “verification mode request”). For example, the resource interface **828** (FIG. 8) may submit the request to the policy enforcement component **830** as usual, but inhibit any actions that would usually result from a “permit” evaluation result. Regardless of the evaluation result, the resource interface **828** may respond to the request with a verification mode report. For example, the resource interface **828** and/or a controlling workflow may invoke an interface element of the verification mode reporting interface **712** (FIG. 7) to generate the verification mode report. The use of dedicated interface elements **914-922** is only one technique for indicating that a request is to be processed in verification mode. Alternate embodiments may utilize any suitable technique. For example, interface elements **904-912** may incorporate an explicit indicator (e.g., a “flag”) for indicating the request is a verification mode request. The resource interface **902** may include an explicit set of interface elements (not shown in FIG. 9) for activating and deactivating verification mode processing. As a further alternative, or in addition, receipt of a verification mode token (e.g., issued by the verification mode service **702**) may act as an indication that an associated request is to be processed in verification mode.

Each interface element **904-922** may correspond to a request (e.g., the request **544** of FIG. 5) subject to a subset of the policies **534-542, 546**. In order to evaluate the request with respect to the subset of the policies **534-542, 546**, the resource interface **828** (FIG. 8) receiving the request may submit the request to the policy enforcement component **830**. FIG. 10 depicts aspects of an example policy enforcement component **1002** in accordance with at least one embodiment. The policy enforcement component **1002** of FIG. 10 is an example of the policy enforcement component **830** of FIG. 8.

The policy enforcement component **1002** may include a policy collector **1004** configured at least to collect relevant policies **534-542, 546** (FIG. 5) from locations throughout the virtual resource provider **502**, and store them local to a decision engine **1006** in a policy cache **1008**. Collected policies may not be in normal form **402** (FIG. 4). The policy enforcement component **1002** may further include a policy normalization component **1010** having functionality corresponding to that of the policy normalization component **610** (FIG. 6) of the policy management service **602**. The policy normalization component **1010** may normalize the set of

policies in the policy cache **1008**. The decision engine **1006** may evaluate requests submitted to the policy enforcement component **1002** with respect to relevant policies in the policy cache **1008**. At times, additional data may be required to support particular decisions with respect to particular policies. The policy enforcement component **1002** may still further include one or more decision data collectors **1012** configured at least to collect the required decision support data ("decision data") from locations throughout the virtual resource provider **502**. Collected decision data may be stored local to the decision engine **1006** in a decision data cache **1014**.

The policy collector **1004** may update the policy cache **1008** responsive to policy update notifications, for example, received from the policy management service **602** (FIG. 6) and/or the virtual resource servers **804-806** (FIG. 8). The policy collector **1004** may subscribe to notifications of updates to relevant policies and/or policy sets maintained at the policy management service **602** and/or the virtual resource servers **804-806**. Alternatively, or in addition, the policy collector **1004** may periodically search for changes to policies in the policy cache **1008** and/or for newly relevant policies.

The decision data collector(s) **1012** may include one or more data collectors configured to interact with the decision data providers **512** (FIG. 5). The decision data collector(s) **1012** may collect decision data responsive to requests by the decision engine **1006**, for example, requests by the decision engine **1006** for decision data that is not present in the decision data cache **1014** and/or that is out-of-date. The decision data collector(s) **1012** may furthermore maintain one or more subsets and/or types of decision data in the decision data cache **1014**, for example, responsive to update notifications from the decision data providers **512** and/or with updates discovered by periodic searching.

As described above with reference to FIG. 3, the control plane **302** may be facilitated by one or more workflows maintained by the workflow management component **322**. FIG. 11 depicts example steps for workflow management in accordance with at least one embodiment. At step **1102**, a request may be received by an interface of the control plane **302** (FIG. 3). For example, one of the user interfaces **304** of the control plane **302** may receive the request from a user, customer and/or administrator of the virtual resource provider **202**. The request may be a call to a programming interface such as an application programming interface (API) or a Web services interface, e.g., utilizing a Simple Object Access Protocol (SOAP) call. Alternatively, or in addition, the request may be received through a graphical user interface (GUI) such as a Web-based GUI. At step **1104**, the request may be analyzed to determine one or more actions required to successfully complete the request. For example, the provisioning interface **310** may analyze the request, and determine a set of actions required to provision a set of virtual resources **218-220** (FIG. 2). When an interface element receiving the request corresponds to a specific action to be performed, the interface may extract information from the request to be utilized in determining aspects and/or parameters of the action to be performed.

At step **1106**, a request may be sent to create a workflow based at least in part on the one or more actions determined at step **1104**. For example, provisioning interface **310** (FIG. 3) may send the request to the workflow management component **322**. The request to create the workflow may include the action(s), action metadata such as type of action, and/or action parameters. In at least one embodiment, the control plane **302** and/or the workflow management com-

ponent **322** maintains a job queue for such requests, and workflows are created responsive to new additions to the job queue. At step **1108**, a workflow and one or more component tasks may be created. For example, the workflow management component **322** may analyze the request of step **1106** to determine the appropriate workflow and component tasks to create.

At step **1110**, execution of the component task(s) may be guided in accordance with the workflow. For example, the workflow management component **322** (FIG. 3) may activate elements of interfaces of various implementation resources to provision the set of virtual resources. Alternatively, or in addition, the workflow management component **322** may manage bids for execution of the component task(s) by components of the virtual resource provider **202**. At step **1112**, it may be determined whether the workflow has finished. For example, the workflow management component **322** may determine whether a final task in a sequence of tasks managed by the workflow has completed. If so, a procedure incorporating step **1112** may progress to step **1114**. Otherwise the procedure may return to step **1110** for a next task and/or task sequence. Workflows may guide multiple task sequences executing in parallel. In this case, it may be that the workflow is not finished until each of the multiple task sequences completes and/or an explicit workflow finished flag is set by one of the component tasks. At step **1114**, the sender of the request of step **1102** may be informed of result(s) of the action(s).

FIG. 12 shows an illustrative example of a process **1200** which may be used to provide computing resources in accordance with an embodiment. Some or all of the process **1200** (or any other processes described herein, or variations and/or combinations thereof) may be performed under the control of one or more computer systems configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors.

In an embodiment, the process **1200** includes receiving **1202** a request for a set of virtual machine instances. The request for the set of virtual machine instances may be made, for instance, by a customer of a remote computing services provider. The request for a set of virtual machine instances may specify one or more actions to be taken in connection with the requested set of virtual machine instances, such as provisioning the virtual machine instances. Other information may also be received as part of or, generally, in connection with the request, such as one or more machine images to be used with the set of virtual machine instances and other configuration details. However, the request may be any general request that identifies a set of virtual machine instances.

In an embodiment, authorization specifications for the requested virtual machine instances are received **1204**. The authorization specifications may be received separately from, or as a part of, the request for virtual machine instances. In embodiments, the authorization specifications include information specifying a policy and/or granularity information. Information specified in the policy may be in accordance with the above description, although any method of specifying actions that may or may not be taken by one or more virtual machine instances, possibly under one or more conditions, may be used. As an example, the policy

21

may specify that all users may use the virtual machine instances to read a particular storage volume that is maintained by a block data storage service. Similarly, another policy may specify that only certain users may use the identified virtual machine instances to write to the storage volume. Generally, the policy may specify one or more actions that may be taken by one or more virtual machine instances. The actions may specify, for example, application programming interface (API) calls that may be made. The API calls may be selected from a plurality of possible calls that may be made in connection with various services provided by a remote computing services provider, or generally, any operator of computing services. As a specific example, a policy may state that any virtual machine instance having particular credentials may make certain API calls that can terminate, stop, and/or otherwise affect other virtual machine instances.

The granularity information, in an embodiment, is information which may be used to determine how credentials for the requested virtual machine instances are to be distributed. For a set of one or more virtual machine instances, in an embodiment, the granularity information is determinative of how many sets of credentials are to be generated for the set of virtual machine instances and is determinative of one or more subsets of the set of virtual machine instances to which the credentials are to be distributed. As an example, the granularity information may specify that each of the requested virtual machine instances should receive a corresponding unique set of credentials. The granularity information may also specify other levels of granularity. For instance, granularity information may specify that one or more groups of the requested virtual machine instances should receive a corresponding set of credentials, for some definition of the groups, which may be user-defined and/or selected. For instance, the granularity information may identify a plurality of subsets of the set of virtual machine instances where each subset of the set of requested virtual machine instances is to receive a separate set of credentials. The subsets may be identified based at least in part on one or more characteristics of the subsets and generally the parameters for the granularity information may include information identifying a host class, a placement group or other logical grouping of computing resources, a geographic location, one or more hardware characteristics for hardware used to provision the virtual machines, and generally any information which may differentiate virtual machine instances from one another, including, but not limited to, user-assigned tags for the virtual machine instances. Users may also define subsets by selecting virtual machine instances for each subset.

In an embodiment, the set of virtual machine instances is provisioned **1206**. Provisioning may be accomplished by allocating hardware resources to the requested virtual machine instances, as described above. In addition, in an embodiment, credentials for the virtual machine instances are generated **1208** in accordance with an embodiment. The credentials may be generated according to the authorization specifications. Continuing with an example provided above, if the authorization specifications indicate that each of the virtual machine instances should have their own credentials, credentials may be generated for each of the virtual machine instances. Generally, if the authorization specifications indicate that each of a plurality of subsets of the virtual machine instances should receive their own credentials, credentials may be generated for each subset. It should be noted that, while generating credentials for the virtual machine instances is shown as a set occurring subsequent to other

22

steps in the process **1200**, credential generation may occur at other times. For instance, credentials may be pre-generated and assigned as needed. In this example, generating credentials may comprise assigning already-generated credentials as necessary, based at least in part on granularity information.

Once the credentials for the virtual machine instances are generated **1208**, the generated credentials are distributed **1210** to the virtual machine instances. Distribution of the credentials for virtual machine instances may be accomplished in any suitable manner, such as by sending, over a communications network, the credentials to the virtual machine instances for use by the virtual machine instances in authenticating themselves with respect to one or more virtual resource services, such as in a manner described above. The distributed credentials may be made available to one or more processes executing on the virtual machine instances, such as by configuring a virtual machine instance that has received the credentials to be able to provide the credentials to the one or more processes and/or storing the distributed credentials in a data store accessible to the virtual machine instance. As another example, the credentials may be distributed and/or otherwise made available via an unauthenticated web server located at a well-known network location, such as a predetermined IP address for a server configured to provide credentials to virtual machine instances and/or other computing resources. Credentials may also be made available via a public-key cryptography system (PKCS) interface. In an embodiment, the public-key cryptography system interface is an interface that complies with PKCS#11 and/or other suitable standards. The distributed credentials may be maintained in a manner wherein the distributed credentials are accessible only by the virtual machine instances to which the distributed credentials were distributed, possibly subject to one or more exceptions. For example, the distributed credentials may be accessible by a computer system involved in the management and distribution. Thus, another virtual machine instance to which the credentials have not been distributed may be prevented from accessing the credentials unless another machine instance to which the credentials were distributed has provided the credentials to the other virtual machine instance and any applicable policies allow such distribution. Distribution of the generated credentials may be done in a secure manner, such as over a secure shell (SSH) connection and, generally, may be done in any suitable manner which results in credentials being usable by virtual machine instances for authentication. Distribution of credentials may also include updating one or more systems used for authentication, such as an authentication service described above, which may maintain a mapping of credentials to virtual machine instances, such as by one or more database tables. In an embodiment, one or more data structures (such as a relational database table) are used to associate, either directly or indirectly, credentials with one or more actions that may be taken by virtual machine instances (or other resources) that are able to provide the credentials, possibly subject to one or more other conditions. An authentication service may utilize the one or more data structures in one or more determinations whether to provide access to a virtual machine instance requesting the access. Upon distribution of credentials, such data structures may be updated.

In an embodiment, one or more policies applicable to the authorization specifications are updated **1212**. For instance, if the authorization specification specifies that one or more subsets of the requested virtual machine instances will have certain privileges, one or more stored policies may be

23

updated to reflect the privileges granted to the requested virtual machine instances. As noted, one or more policies may be based at least in part on credentials. As an example, a policy may permit or deny one or more actions to virtual machine instances having a particular set of credentials and/or a certain class of credentials. Accordingly, updating policies may include identifying in the policies information that allows specific credential-based conditions to be checked. Updating policies may also include, for one or more policies associated with the credentials, assigning one or more virtual machine instances to which credentials have been distributed as principals of the one or more policies. As with other steps in the process 1200 (and generally, with steps of other processes described herein), updating policies is shown in a particular place in connection with a series of steps. However, policy updates may be made at any suitable time.

Additional actions may be taken in connection with the process 1200, or variations thereof. For example, the credentials may be received from the virtual computing instances and used in an authentication process, as described below. In addition, more than one set of credentials may be provided to one or more virtual machine instances. Further, while various embodiments described herein focus on authentication in connection with services of a remote computing services provider, credentials distributed in accordance with the various embodiments may be used in other instances, such as for authentication with third parties and, generally, any process in which credentials are used.

As noted above, credentials may be generated in accordance with information specifying a granularity for such credentials. Accordingly, FIG. 13 shows an illustrative example of the process 1300 which may be used to generate and distribute credentials according to granularity information specified by the user. In an embodiment, a granular unit is identified 1302. A granular unit may be a set of one or more virtual machine instances associated with a set of credentials used by all of the virtual machine instances in the set. The granular unit may be based in part on granularity information provided by a user in connection with the request for one or more virtual machine instances. For example, if the granularity information indicates that each virtual machine instance is to receive its own set of credentials, a granular unit may be a virtual machine instance. Granular units may also be sets of multiple virtual machine instances, such as indicated by the granularity information. A granular unit may be, for example, a set of all virtual machine instances of a customer of a remote computing services provider that are implemented using a particular class of host hardware, or a subset of such virtual machine instances. As another example, a granular unit may be a set of a customer's (or, generally, user's) virtual machine instances located in a particular data center, or having been assigned with a tag. As yet another example, a granular unit may be a set of a user's virtual machine instances having a particular property or combination of properties, or, generally, satisfying a set of one or more criteria.

In an embodiment, credentials for the identified granular unit are generated 1304 in accordance with an embodiment. Generation of the credentials may be performed in any suitable manner, which may depend on one or more protocols for which the credentials will be used. In an embodiment, credentials are generated as unique identifiers for corresponding sets of virtual machine instances. Credentials may be or include a key for an authentication algorithm, such as a key for a public-key cryptography algorithm. Thus, for an identified granular unit, the credentials are a unique

24

identifier string to be shared by the virtual machine instances of the identified granular unit. However, the credentials may be any set of credentials that may be used with an authentication protocol. In addition, the credentials may encode at least some information representative of metadata of resources in the granular unit. For instance, if the granular unit is a virtual machine, the credentials may encode an Internet Protocol (IP) address for the virtual machine instance, and, if the identified granular unit includes a plurality of virtual machine instances, the credentials may encode a set of IP addresses for the virtual machine instances of the identified granular unit. Generally, the credentials may encode any information about any virtual machine instances associated with the credentials. In an embodiment, credentials encode one or more policies. For instance, if a virtual machine instance is implemented on a host machine, credentials accessible by the virtual machine may encode a token that may be required to be used with a key provided by the host machine in an authentication procedure. In this manner, the credentials encode a policy that requires use of the credentials to be limited to virtual machines implemented by the host. Such credentials may be provided to multiple virtual machines on the same host. Thus, if another virtual machine provides the credentials from another host that does not have the corresponding key, authentication may fail. In addition to the above, as noted, credentials, or portions thereof, may be pre-generated and assigned as needed instead of generated upon identification of a granular unit as illustrated in FIG. 13.

In an embodiment, the generated credentials are added 1306 to a credentials mapping. As discussed, the mapping may associate credentials of virtual machine instances with the instances themselves. For instance, the mapping may associate credentials to a set of IP addresses and/or other identifiers of virtual machine instances. As noted above, the mapping may be maintained by an authentication service. For a virtual machine, the mapping may be used, in accordance with the above description, to verify, for instance, that a request made by a virtual machine using the credentials is associated with the credentials, to ensure that a virtual machine has the proper privileges to perform any actions requested by the virtual machine, and/or to allow the virtual machine to establish a connection (such as an SSH, HTTP, or other connection) with another system, such as a system providing a virtual resource service, as described above.

In an embodiment, a determination is made 1308 whether there are additional granular units for which credentials could be generated. If, for example, credentials have not been generated to all granular units specified by authentication information specified by a customer, the determination may be that there are additional granular units. In an embodiment, if the determination is that there are additional granular units for which credentials should be generated, the next granular unit is identified 1310 in accordance with an embodiment. The next granular unit may be the same type as the previously-identified granular unit, although it may be of a different type. For instance, if the previously-identified granular unit was a set of one or more virtual machines, another identified granular unit may have more or less members of a set of one or more virtual machines. With the next granular unit identified 1310, credentials for the identified granular unit are generated 1304, such as in a manner described above. If, however, it is determined that there are no additional granular units, then the generated credentials are distributed 1312 to the computing resources of the granular unit or units for which credentials were generated. As an example, if credentials were generated for each of a

plurality of virtual machine instances, the generated credentials would be distributed to each of the virtual machine instances. Similarly, if credentials were generated for a granular unit having multiple virtual machine instances, the credentials generated for the granular unit may be distributed to the virtual machines of the granular unit. In this last example, each virtual machine instance in the granular unit may receive the same credentials (such as the same unique identifier), while virtual machines may receive other credentials.

As discussed, credentials distributed to virtual machine instances in accordance with various embodiments described herein can be used in various ways. FIG. 14 shows an illustrative example of one such process 1400 which shows how credentials generated in accordance with the various embodiments may be used. In an embodiment, the process 1400 includes receiving 1402 credentials in connection with a request to access a computing resource. A request to access a computing resource may be received, such as in the manner described above. For instance, a resource interface of a virtual resource service may receive a request generated from another virtual resource service that operates virtual machine instances on behalf of one or more customers and/or other users. As a concrete example, a resource service for virtual machine instances may send a request originating from one of the virtual machine instances to a block storage or other storage resource. In an embodiment, the request includes the credentials from a virtual machine instance that made the request, although the credentials may be provided separately from the request, such as before the request, or after a subsequent authentication request for credentials from the receiver of the request. Generally, the credentials may be received in any suitable manner, and the manner by which credentials are received may vary according to various protocols being used in accordance with the various embodiments.

When the credentials in connection with the request to access the computing resource are received 1402, in an embodiment, a determination is made 1404 whether the requester has a privilege allowing the requested access. Determination may be made, for instance, by a policy enforcement component of a virtual resource service such as in the manner described above in connection with FIG. 8. The policy enforcement component may utilize an authentication service to authenticate the credentials and, if successfully authenticated, establish a connection with the virtual machine instance that provided the credentials, at least for a period of time. In some embodiments, however, the receiver of the credentials may authenticate a virtual machine supplying the credentials itself instead of using a separate authentication service. Generally, any process for determining whether the credentials received are valid may be used.

If it is determined that the requester has a privilege allowing the requested access, the requested access may be provided 1406. Providing the requested access may be performed in any suitable manner. In an embodiment, providing the requested access includes establishing an SSH or other connection with a computing device of a resource, such as a server operating in connection with a virtual computing resource. As another example, the virtual machine instance requesting the access may be provided read, write, and/or other access to a data store of a data storage service. As another example, a virtual machine instance running one or more applications (such as Red Hat Enterprise Linux (RHEL)) may be provided access to an update service for the one or more applications, such as an

RHEL update repository. In this manner, the virtual machine instance can use its credentials to update applications running on the virtual machine instance. It should be noted that, credentials may be used in different ways in different circumstances. For instance, credentials at some times may be used to establish a connection with another computing device while at other times credentials may be used to simply cause performance of one or more operations at another computing resource.

If it is determined that the requester does not have a privilege allowing the requested access, then the requested access is denied 1408. As illustrated in FIG. 14, although not necessary, one or more other actions may be completed in addition to denying the requested access. For instance, as illustrated, one or more users are notified 1410 of a security violation. Users may receive, for example, an electronic message specifying that a security violation has occurred, and including details of the security violation. In addition, as an example, remedial action may be taken 1412 in accordance with an embodiment. As an example, if the request to access a computing resource included a set of credentials and those credentials were used to make the determination that the requestor did not have a privilege allowing the requested access, the credentials may be permanently or temporarily invalidated so that the credentials are not used in attempts to access other computing resources. Generally, any remedial action may be taken. Further, as with other processes described herein, the process 1400 is shown in a simplified manner and variations are contemplated as being within the scope of the present disclosure. For instance, any remedial action taken may be based at least in part on additional information, such as a number of failed authentication attempts, the type of access requested, information encoded in the request, such as an IP address or other information, and the like.

In some instances, it may be desirable to manage the validity of credentials in order to increase security. For instance, in an embodiment, if credentials belong to a single virtual machine instance, and the virtual machine instance currently is not provisioned, but stored in a data store for potential future use, the virtual machine may be unable to use the credentials. Thus, it may be desirable to prevent the credentials from being used when no virtual machine is able to use the credentials to authenticate itself, either because of a state of one or more virtual machines having the credentials and/or because of one or more policies. Accordingly, FIG. 15 shows an illustrative example of a process 1500 for managing credentials, in accordance with an embodiment. In this illustrative example, the process 1500 includes distributing credentials to one or more virtual machine instances. Distribution of the credentials may be performed in any suitable manner such as those described above.

In an embodiment, provisioning of virtual machine instances is detected 1504. Detection of provisioning of the computing resources may be performed in any suitable manner. In an embodiment, provisioning a virtual machine instance is accomplished according to a workflow that includes sending a message that the virtual machine instance has been successfully provisioned. Accordingly, detecting the provisioning of virtual machine instances may be performed by receiving one or more such messages. It should be noted that, while detecting provisioning of virtual machine instances is used to illustrate an embodiment of the present disclosure, other actions may be taken in place of detection of provisioning the virtual machine instances, such as receiving a request to provision the virtual machine

instances, or, generally, receiving any signal indicative of a part of a provisioning process.

In an embodiment, when the provisioning of the virtual machine instances is detected **1504** (or, generally, when an event in connection with the provisioning of the virtual machine instances is detected), the credentials are activated for the virtual machine instances. Activating the credentials may be performed, for instance, by updating a credentials mapping of credentials to virtual machine instances to reflect the credentials being active. Such may be accomplished by adding the credentials into the mapping, by tagging the credentials in a mapping as active, or, generally, in any suitable manner. As a concrete example, a database table or other mechanism of an authentication service that associates credentials with virtual machine instances may be updated to reflect whether the credentials are currently active by adding an appropriate number of rows into the table or setting a value for the credentials to indicate that the credentials are active. If credentials are active, then virtual machine instances having the credentials may use the credentials for authentication. If the credentials are inactive, then any device attempting to use the credentials for authentication should be unable to authenticate. Generally activating the credentials for the provision of computing resources can be done in any manner that allows the provisioned computing resources to utilize the credentials for authentication with one or more other virtual resource services. For instance, if a virtual machine wants to access a data store of a storage service, the virtual machine may use the credentials to authenticate itself with the storage service because the credentials are active. The credentials may remain active as long as the computing resources remain provisioned.

In an embodiment, deprovisioning of one or more of the provisioned virtual machine instances is detected **1508**. While deprovisioning is used for the purpose of illustration, any suitable action may be detected, including, but not limited to, actions related to, but not necessarily equivalent to the deprovisioning of provisioned virtual machine instances. An example includes receiving a command to deprovision one or more of the provisioned virtual machine instances. Returning to the illustrated example, detection of deprovisioning of the one or more provisioned virtual machine instances may be made when hardware for a virtual machine instance is deallocated from the virtual machine instance and a memory state of the virtual machine instance is stored for future use.

In an embodiment, upon detection of the deprovisioning of one or more of the deprovisioned virtual machine instances, the credentials for the deprovisioned virtual machine instances are deactivated **1510** in order to ensure that credentials cannot be used by any other computing resources. In this manner, if a security breach occurs with respect to a deprovisioned virtual machine, the credentials cannot be used to gain unauthorized access to one or more computing resources. For example, if all of the virtual machine instances sharing the same credentials are deprovisioned, then the credentials would not be active for any virtual machine and, therefore, cannot be used for unauthorized authentication. Similarly, in embodiments where virtual machines authenticate themselves using the credentials and information about themselves (such as an IP address) for authentication, authentication using the credentials may be impossible or, at least, more difficult. In one embodiment, if less than all of the provisioned virtual machine instances sharing the same credentials are deprovisioned, the credentials may be deactivated only for the deprovisioned virtual

machine instances, while the remaining provisioned virtual machine instances may continue to use the credentials.

In some instances, it may be desirable to take different actions with respect to credentials when different events occur. For instance, in some instances it may be desirable to invalidate credentials upon detection of a security breach while in other instances it may be desirable to keep credentials valid while taking other remedial action. Accordingly, FIG. 16 illustrates an illustrative example of a process for managing credentials in accordance with an embodiment. In an embodiment, a security breach of a virtual machine instance is detected **1602**. A detection of a security breach may be done in any suitable manner or combination thereof. For instance, detection of the security breach may be made upon a determination that credentials used in a request to access a computing resource were unsuccessfully used. As another example, the security breach may be detected based at least in part on access to a virtual machine instance by an unauthorized user. Other ways of detecting the security breach may be used as well, including, but not limited to, manual detection, detection of abnormal operation and/or other automatic methods, and/or generally any suitable method.

In an embodiment, a determination is made **1604** which action to take upon detection of the security breach. The determination may be based, at least in part, on one or more parameters such as a type of breach that is detected. While, for the purpose of illustration, three different actions are shown, there may be greater than or fewer than three actions, and the actions may be different than those illustrated. In an embodiment, upon detection of the security breach, one of several actions are selected. The selection may be made from actions such as invalidating the credentials used when the security breach was detected, modifying an access policy associated with a virtual machine instance involved with the security breach, or rotating credentials involved with the security breach. If the determination is made **1604** to invalidate the credentials, then the credentials are invalidated **1606**. For instance, credentials may be deleted from a mapping that associates credentials with virtual machine instances, may be marked as inactive, or otherwise invalidated. In an embodiment, if the determination is to modify an access policy, then the policy is modified **1608** in order to revoke privileges from the instance involved in the security breach. Using the illustrated policies described above as an example, a condition(s) component of a policy may be modified to revoke one or more privileges from one or more virtual machine instances. An action(s) component or other applicable component of a policy may also be updated. Generally, the policy may be updated in any suitable manner.

If the determination is to rotate the credentials, then new credentials are generated **1610** in accordance with an embodiment, although, as discussed, credentials may have been pregenerated and stored for use when needed. In an embodiment, when new credentials are generated, a determination is made **1612** whether to distribute the new generated credentials to the compromised instance. The determination may be made based at least in part on the type of security breach that was detected. For instance, if it would present an undue risk to provide new credentials to a compromised virtual machine instance, the determination may be to not distribute the new generated credentials to the compromised instance. An example of such an undue risk includes an instance that was compromised in such a way that one or more unauthorized applications have been loaded on the compromised instance. If it would not create an undue

risk to provide new credentials to the compromised virtual machine instance, then the determination may be to provide new generated credentials to the compromised virtual machine instance. If it is determined to distribute the new credentials to the compromised instance, then the generated credentials are distributed **1604** to all instances in a granular unit of the compromised instance. For instance, the new credentials may be provided to all virtual computing instances that shared the credentials involved in the detected security breach.

If, however, a determination is not to distribute the new credentials to the compromised instance, then the credentials are distributed **1606** to all instances in a granular unit of the compromised instance except for the compromised instance. For example, the new credentials may be provided to all virtual machine instances sharing the credentials involved in the detected security breach, except for the compromised virtual machine instance. If the granular unit only contains the compromised instance, then the generated credentials may not be distributed at all and, in some embodiments, may not have been generated.

While not illustrated in the figure, the generation and distribution of new credentials may also include invalidation of previous credentials, as described above. In addition, as mentioned, other variations are also contemplated as being within the spirit of the present disclosure. For instance, actions different from those discussed above may be taken in response to detection of a security breach. In addition, combinations of actions may be taken in response to detection of a security breach. For instance, credentials may be rotated, one or more policies may be modified, and credentials involved in the security breach may be invalidated all in response to the security breach being detected.

Other variations are also contemplated as being within the spirit of the present disclosure. For instance, for the purpose of illustration, various embodiments are discussed in connection with virtual machine instances. However, various other embodiments may involve other virtual computing resources. For instance, any virtual computing resource that uses credentials to access another virtual computing resource may utilize various embodiments described herein, and variations thereof. Generally, techniques described and suggested herein are not limited to virtual computing instances. In addition, various techniques described and suggested herein may be used in connection with virtual computing resources in addition to virtual machine instances.

Generally, the various embodiments described herein may be implemented in a wide variety of operating environments, which in some cases may include one or more user computers, computing devices, or processing devices which may be utilized to operate any of a number of applications. User or client devices may include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless, and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also may include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices also may include other electronic devices, such as dummy terminals, thin-clients, gaming systems, and other devices capable of communicating via a network.

Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting com-

munications using any of a variety of commercially-available protocols, such as TCP/IP, OSI, FTP, UPnP, NFS, CIFS, and AppleTalk. Such a network may include, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network, and any combination thereof. The network may, furthermore, incorporate any suitable network topology. Examples of suitable network topologies include, but are not limited to, simple point-to-point, star topology, self organizing peer-to-peer topologies, and combinations thereof.

In embodiments utilizing a Web server, the Web server may run any of a variety of server or mid-tier applications, including HTTP servers, FTP servers, CGI servers, data servers, Java servers, and business application servers. The server(s) also may be capable of executing programs or scripts in response to requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java®, C, C# or C++, or any scripting language, such as Perl, Python, or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle®, Microsoft®, Sybase®, and IBM®.

The environment may include a variety of data stores and other memory and storage media as discussed above. These may reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (“SAN”) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers, or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device may include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (e.g., a mouse, keyboard, controller, touch screen, or keypad), and at least one output device (e.g., a display device, printer, or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random access memory (“RAM”) or read-only memory (“ROM”), as well as removable media devices, memory cards, flash cards, etc.

Such devices also may include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader may be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules including program modules, services, or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be utilized and/or particular elements might be implemented in hardware, software (including portable soft-

31

ware, such as applets), or both. Further, connection to other computing devices such as network input/output devices may be employed.

Storage media and computer readable media for containing code, or portions of code, may include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules, or other data, including RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which may be utilized to store the desired information and which may be accessed by the a system device. Program modules, program components and/or programmatic objects may include computer-readable and/or computer-executable instructions of and/or corresponding to any suitable computer programming language. In at least one embodiment, each computer-readable medium may be tangible. In at least one embodiment, each computer-readable medium may be non-transitory in time. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

The use of the terms “a” and “an” and “the” and similar referents in the context of describing embodiments (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. The term “connected” is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments and does not pose a limitation on the scope unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of at least one embodiment.

Preferred embodiments are described herein, including the best mode known to the inventors. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate, and the inventors intend for embodiments to be constructed otherwise than as specifically described herein. Accordingly, suitable embodiments include all modifications and equivalents of the subject

32

matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is contemplated as being incorporated into some suitable embodiment unless otherwise indicated herein or otherwise clearly contradicted by context.

All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

What is claimed is:

1. A computer-implemented method, comprising:
 - under the control of one or more computer systems, configured with executable instructions, that cause the method to be performed by one or more hardware processors,
 - receiving, for a set of virtual computing resources that are collectively under common control, an instruction to distribute credentials for at least a subset of two or more of the set of virtual computing resources, the instruction including granularity information usable to:
 - determine how many or the credentials are to be generated; and
 - determine at least the subset; and
 - making available, to the at least a subset, corresponding credentials, the corresponding credentials usable by the at least a subset to authenticate with one or more systems managing one or more other computing resources thereby enabling the at least a subset to access the one or more other computing resources, the corresponding credentials made available to the at least a subset being inaccessible to at least the other virtual computing resources.
2. The computer-implemented method of claim 1, further comprising generating the corresponding credentials in response to receipt of the instruction.
3. The computer-implemented method of claim 1, further comprising:
 - generating at least one policy associated with the credentials made available, the policy identifying a set of one or more virtual computing resources to which the corresponding credentials have been made available; and
 - regulating, based at least in part on the policy, access to the one or more computing resources.
4. The computer-implemented method of claim 1, wherein the virtual computing resources include virtual machine instances and wherein the at least a subset includes a plurality of the virtual machine instances.
5. The computer-implemented method of claim 1, wherein making available corresponding credentials includes providing the corresponding credentials from a web server having a predetermined network location.
6. The computer-implemented method of claim 1, wherein the corresponding credentials encode one or more policies.
7. The computer-implemented method of claim 1, wherein the virtual computing resources are virtual machine instances.
8. The computer-implemented method of claim 1, wherein the granularity information usable to determine at least the subset includes information specifying one or more characteristics for distinguishing the at least a subset from other virtual computing resources of the set of virtual computing resources that are under the common control.

33

9. The computer-implemented method of claim 1, wherein the at least a subset contains multiple virtual machine instances.

10. The computer-implemented method of claim 1, wherein:

the at least a subset includes a plurality of specified subsets; and

the method further comprises:

detecting a change of state of a particular subset of the plurality of specified subsets; and

in response to detecting the change of state, causing credentials corresponding to the particular subset to be at least temporarily unusable.

11. One or more non-transitory computer-readable storage media having collectively stored thereon instructions executable by one or more processors of a computer system that, when executed by the one or more processors, cause the computer system to at least:

receive, for a set of virtual computing resources that are collectively under common control, granularity information determinative of at least one subset of two or more of the set of virtual computing resources and determinative of how many credentials are to be distributed to the at least one subset; and

distribute, to the at least one subset, the credentials for use by the at least one subset in authenticating with one or more systems managing access to one or more other computing resources, thereby enabling the set of virtual computing resources to access one or more other computing resources.

12. The one or more non-transitory computer-readable storage media of claim 11, wherein the set of virtual computing resources include virtual machine instances.

13. The one or more non-transitory computer-readable storage media of claim 12, wherein the at least one subset is at least one set of two or more virtual machine instances.

14. The one or more non-transitory computer-readable storage media of claim 11, wherein the granularity information further identifies the at least one subset by one or more characteristics that distinguish the at least one subset from other virtual computing resources of the set of virtual computing resources.

15. The one or more non-transitory computer-readable storage media of claim 14, wherein the instructions further cause the computer system to:

34

detect a change of state of the at least one subset; and in response to detecting the change of state, cause credentials corresponding to the at least one subset to be at least temporarily unusable.

16. A computer system for managing credentials, comprising:

one or more hardware processors; and

memory, including executable instructions that, when executed by the one or more processors, cause the one or more hardware processors to collectively at least:

distribute, to each subset of a plurality of subsets of a set of computing resources that are collectively under common control, different credentials usable to authenticate with a system managing access to one or more other computing resources, each subset including two or more computing resources of the set of computing resources that are identifiable based at least in part on one or more characteristics that distinguish the two or more computing resources of the subset from computing resources of each other subset of the set of computing resources;

detect an event in connection with at least one computing resource of a subset of the plurality of subsets, the at least one computing resource being associated with corresponding credentials of the distributed different credentials, the corresponding credentials being inaccessible to at least computing resources under the common control that are outside of the subset; and

in response to detection of the event, take an action that modifies an ability of the at least one computing resource to use the corresponding credentials.

17. The computer system of claim 16, wherein the event includes a security breach of the computing resource for which the event was detected.

18. The computer system of claim 17, wherein the action includes invalidating the corresponding credentials.

19. The computer system of claim 16, wherein the event includes a change to an inactive state of the computing resource for which the event was detected.

20. The computer system of claim 19, wherein the action includes causing the corresponding credentials to become at least temporarily inactive.

* * * * *